IBM FlashSystem A9000R

© Copyright IBM Corp. 2025.
US Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Tables of Contents

IBM FlashSystem A9000R welcome page 1

Cross-system feature use in the IBM Spectrum Accelerate family 1

Platform and application integration 4

GDPR and IBM systems 5

What's new 5

Agreements and license information 8

Warranty information 8

Notices 8

- Trademarks 9
- Homologation statement 9
- Safety and environmental notices 9
  - Safety notices and labels 9
  - Danger notices for IBM FlashSystem A9000R systems 10
  - Caution notices for IBM FlashSystem A9000R systems 11
  - Special caution and safety notices 11
    - Laser safety 12
    - Ladder usage 12
    - Fire suppression systems 12
    - Sound pressure 12
    - Leakage current 12
    - Site preparation 13
  - Environmental notices 13
- Electromagnetic compatibility Class A notices 14
  - Canada Notice 14
  - European Community and Morocco Notice 14
  - Germany Notice 14
  - Japan Electronics and Information Technology Industries Association (JEITA) Notice 15
  - Japan Voluntary Control Council for Interference (VCCI) Notice 15
  - Korea Notice 15
  - People's Republic of China Notice 16
  - Russia Notice 16
  - Taiwan Notice 16
  - United Kingdom Notice 16
  - United States Federal Communications Commission (FCC) Notice 16
- Electromagnetic compatibility Class B notices 17
  - Canada Notice 17
  - European Community and Morocco Notice 17
  - Germany Notice 17
  - Japan Electronics and Information Technology Industries Association (JEITA) Notice 18
  - Japan Voluntary Control Council for Interference (VCCI) Notice 18
  - Taiwan Notice 18
  - United States Federal Communications Commission (FCC) Notice 19

Release notes 19

- Release notes for IBM FlashSystem A9000R 19
  - What's new in 12.3.2.e 19
  - Hot upgrade from earlier versions 20
  - Management and solutions 20
  - Change log 21
    - 12.3.2.e (January 2023) 22
    - 12.3.2.d (March 2022) 23
    - 12.3.2.c (January 2021) 24
    - 12.3.2.b (May 2020) 25
    - 12.3.2.a-3 (May 2020) 26
    - 12.3.2.a-2 (December 2019) 27
    - 12.3.2.a-1 (September 2019) 27
    - 12.3.2.a (August 2019) 27
    - 12.3.2 (May 2019) 28
    - 12.3.1 (January 2019) 29
    - 12.3.0.a (November 2018) 30
    - 12.3.0 (November 2018) 30
  - Fixes and features of earlier versions 32
  - Limitations 32
  - Known issues 33
  - Known interoperability issues 35
- Release notes for IBM HyperSwap Quorum Witness 36

Release scope and highlights 36
Compatibility and requirements 36
Supported Linux versions 36
Supported storage systems 36
Host requirements 36
Change log 37
1.0.3 (February 2021) 37
1.0.2 (August 2019) 37
1.0.1 (June 2018) 37
1.0.0 (June 2017) 38
Limitations 38
Known issues 38
Release notes for IBM Hyper-Scale Manager 38
Release notes for IBM XCLI Utility 39
Overview 39
What's new in 5.5.4.1 39
Compatibility 39
Change log 39
5.5.4.1 (June 2022) 40
5.5.4 (May 2022) 40
5.2.0.a (July 2017) 40
5.2.0 (April 2017) 40
5.1.0 (November 2016) 40
Limitations 41
Known issues 41
Release notes for IBM Remote Support Proxy 41
What's new in 1.3.2.2 41
Supported operating systems 41
Supported storage systems 41
Limitations 42
Known issues 42
Release note PDF publications 42

Product overview 43
Introduction 43
Architecture 44
Model 415 45
Models 425 and U25 47
Flash enclosures 49
Grid controllers 50
Back-end interconnect 51
Logical architecture 51
Scale-out grid architecture 51
Functionality 51
Specifications 52
Performance 52
Flash-optimized data reduction 52
Data reduction stages 53
Pattern detection and removal 54
Inline deduplication 55
Inline compression 56
Flash-optimized data protection 56
Two-dimensional flash RAID 57
Vaulting mechanism 58
Scrubbing mechanism 59
Flash-optimized data path 59
Capacity management 61
Thin provisioning 61
Allocation limit 62
Deduplication 62
Volumes and snapshots 62
Volume function and lifecycle 62
Snapshot function and lifecycle 64
Creating a snapshot 65
Locking and unlocking a snapshot 66
Duplicating a snapshot 66
Creating a snapshot of a snapshot 66
Formatting a snapshot or a snapshot group 66
Additional snapshot attributes 67
Redirect-on-Write (ROW) 68
Full volume copy 69
Restoring volumes and snapshots 70
Storage pools 72
Consistency groups 73

Snapshot of a consistency group 74
Consistency group snapshot lifecycle 75
Quality of Service (QoS) performance classes 77
Connectivity 77
Physical connectivity 77
iSCSI connectivity 78
Management connectivity 78
Virtual LAN (VLAN) support 79
Host system attachment 79
CHAP authentication of iSCSI hosts 80
Clustering hosts into LUN maps 80
Synchronous remote mirroring 82
Remote mirroring basic concepts 83
Synchronous mirroring operations 84
Synchronous mirroring configuration and activation options 84
Synchronous mirroring statuses 85
Synchronous mirroring role switchover and role change 86
Role switchover when synchronous mirroring is operational 87
Role change when synchronous mirroring is not operational 87
I/O operations in synchronous mirroring 87
Coupling synchronization process 88
Synchronous mirroring of consistency groups 89
Asynchronous remote mirroring 90
Asynchronous remote mirroring 91
Snapshot-based technology in asynchronous mirroring 91
Mirroring-special snapshots 91
Disaster recovery scenarios in asynchronous mirroring 92
High availability with HyperSwap 93
Design guidelines 93
Configuration 93
HyperSwap volumes and consistency groups 94
Quorum Witness 94
Failure domain 94
Topologies 95
Volume level perspective 95
System level perspective 96
Other topologies 97
Automatic recovery scenarios 98
Operating a HyperSwap-based High Availability solution 98
Establishing a HyperSwap relation 99
Monitoring a HyperSwap volume 100
Modifying a HyperSwap volume 101
High Availability snapshots 101
Failure detection and recovery scenarios 101
System A - System B connectivity failure 102
System A - Quorum Witness connectivity failure 102
System A - System B and System A - Quorum Witness connectivity failure or System A failure 103
System B - Quorum Witness connectivity failure 104
System B - System A and System B - Quorum Witness connectivity failure or System B failure 104
System A - Quorum Witness and System B - Quorum Witness connectivity failure 105
System A I/O serving failure 105
Total connectivity failure 106
Multi-site high availability and disaster recovery 106
Multi-site HA/DR advantages 107
Multi-site HA/DR architecture 108
Peers in a Multi-site relation 108
Standby asynchronous mirror 109
Snapshots in Multi-site HA/DR 110
Multi-site relation assembly 110
Asynchronous mirror activation in a multi-site relation 111
Adding a standby asynchronous mirror 111
Multi-site HA/DR operations 111
Monitoring a Multi-site relation 111
Changing roles in a Multi-site relation 112
Deleting a Multi-site relation 112
Downgrading a Multi-site relation to a two-way mirroring relation 113
Recovery scenarios 113
Migrating data 113
Data migration overview 113
I/O handling in data migration 114
Data migration stages 114
Handling failures 116
Volume migration with IBM Hyper-Scale Mobility 116

Volume migration process 117
Mirroring Resources 118
Examples for replication couplings 118
Data-at-rest encryption 119
Encryption key management schemes 119
User roles and permissions 120
User groups 121
Predefined users 121
User information 121
User authentication and access control 122
Native authentication 122
LDAP authentication 122
LDAP authentication logic 123
Multi-tenancy 124
Working with multi-tenancy 126
Management and monitoring 126
Event reporting and handling 127
Event information 127
Event notification rules 128
Event notification destinations 128
Event notification gateways 128
Integration with ISV environments 129
Supporting VMware vStorage extended operations 129
Integration with Microsoft Azure Site Recovery 129
Software upgrade 130
Preparing for software upgrade 130
Remote support and proactive support 130

Deploying 131
Roles and responsibilities 131
Overview 132
Planning best practices and requirements 132
Additional product information 133
System specifications 133
Physical configuration options 135
Components and interconnect 135
Rack configurations 136
FlashSystem A9000R grid elements 137
Model 415 rack configurations 137
Model 425 and U25 rack configurations 138
Conducting system scale-out (MES) 140
Flash enclosure components and feature codes 140
Grid controller components and feature codes 141
Rear-door heat exchanger 142
Height reduced shipping option 143
Weight reduced shipping option 143
Radio frequency identification device option 144
Physical installation site requirements 144
Floor and space requirements 144
Raised or non-raised floor considerations 145
Floor-load requirements 146
Rack dimensions and service clearance requirements 146
Bottom rack dimensions 147
Preparing for raised-floor installation and cabling 147
Preparing for non-raised-floor installation and cabling 148
Preparing for the rear-door heat exchanger 148
Power requirements 148
Power sources 149
Power consumption 150
Input voltages and frequencies 150
Main power cables specifications 150
Main power cables considerations 152
Environmental requirements 153
Operating environment requirements 153
System environmental acclimation 154
Air circulation and cooling 154
Temperature threshold and events 155
Contamination information 155
Acoustic declaration 156
Operating vibration requirements 157
Planning for the rear-door heat exchanger 157
Rear-door heat exchanger operating specifications 158
Rear-door heat exchanger performance 158
Preparing your site to provide water to the rear-door heat exchanger 159

Water specifications for the secondary cooling loop 160
Water delivery specifications for secondary loops 161
Raised floor hose requirements and management 164
Non-raised floor hose requirements and management 166
Secondary cooling loop parts and services information 167
Maintenance schedule 169
Site security considerations 169
Network and host connectivity requirements 169
Network connections for management 170
Management port requirements 170
IP configuration 171
Control (management) protocols 171
Host network connectivity and multipathing configurations 172
Fibre Channel (FC) network configurations 173
Fibre Channel (FC) adapter types 173
Fibre Channel host port configuration 174
Fibre Channel connectivity requirements 174
Fibre Channel best practices 175
Fibre Channel zone types 175
Zone type examples 176
Fibre Channel connectivity during scale-out (MES) 178
Example of adding zone types during scale-out 179
iSCSI network configurations 179
iSCSI host port configuration 179
iSCSI connectivity requirements 180
iSCSI best practices 181
iSCSI connectivity during scale-out (MES) 181
Network cable requirements 181
Network and host connectivity security information 182
Internet Protocol Security (IPSec) 182
Data-at-rest encryption 182
Internal data-at-rest encryption scheme 182
External data-at-rest encryption scheme 182
User authentication and access control 183
Native user authentication 183
External authentication via Lightweight Directory Access Protocol (LDAP) 183
PCI DSS compliance 184
Migration and mirroring connectivity 185
Physical connectivity for mirroring and migration using Fibre Channel 185
Physical connectivity for mirroring and migration using iSCSI 185
Migration and mirroring best practices 186
Planning for physical shipment 187
Planning to receive delivery 187
Planning for relocation 188
Relocation requirements for systems with external encryption schemes 188
Relocation requirements for systems with local encrypted schemes 188
Shipment weights and dimensions 189
Planning for remote support, on-site service, and maintenance 189
Planning for remote support connection 189
Remote support for severe system conditions 190
Planning for Call Home 191
Planning for Call Home Web 191
Required support information 192
Support and software maintenance security information 193
Deploying IBM HyperSwap Quorum Witness 194
Introduction 194
Quorum Witness and storage systems 195
Multi-site HA/DR solution 195
Installation 196
Compatibility and requirements 196
Performing first-time installation 197
Upgrading an existing installation 198
Defining Quorum Witness on IBM storage systems 198
Uninstalling the Quorum Witness software 198
Operation 199
Managing SSL certificates 199
Replacing a self-signed certificate on the Quorum Witness 199
Managing CA public key on the Quorum Witness 199
Checking and controlling the Quorum Witness service 200
Modifying the Quorum Witness communication port 200
Enabling IPv6 connectivity 200

Troubleshooting 201
Checking the log files 201
Best practices 201
Highly-available VM deployment 201
Restoring Quorum Witness database 202
Restoring the database by installing Quorum Witness on a new system 202
Restoring the database by reinstalling Quorum Witness on the existing system 202
Additional reference information 203

Operating 203
Powering on after shutdown 203
Powering off the system 204

Managing 204
Managing systems using IBM Hyper-Scale Manager 204
Managing systems using the command-line interface (CLI) 205

Configuring 205
Configuring host attachment 205
Preparing for host attachment 205
Physically connecting the host to the storage system 206
Physically connecting a host using Fibre Channel (FC) connections 206
Physically connecting a host using iSCSI connections 206
Host connectivity and parameter settings for FlashSystem A9000 and A9000R 206
Setting Fibre Channel connectivity 207
Fibre Channel zoning 208
Setting iSCSI connectivity 209
Configuring iSCSI ports 209
Defining the host on the storage system 210
Monitoring host connectivity 212
Configuring Linux host attachment 212
Red Hat Enterprise Linux (RHEL) 8.4 Fibre Channel host attachment 212
Red Hat Enterprise Linux (RHEL) 8.4 iSCSI host attachment 213
Configuring Windows host attachment 214
Windows Fibre Channel host attachment 214
Windows iSCSI host attachment 215
Configuring Solaris host attachment 215
Configuring key servers 217
Thales CipherTrust Manager k170v v2.3.x and 2.6.x migration to client certificate authentication key 218
Enabling encryption on the storage system 218
Migrating from global key to Client Certificate Authentication 222
CipherTrust Manager cluster configuration 223
Using an external CA certificate 224
Thales CipherTrust Manager k170v v2.2.x migration to client certificate authentication key 224
Enabling encryption on the storage system 225
Migrating from global key to Client Certificate Authentication 228
CipherTrust Manager cluster configuration 229
Using an external CA certificate 230
Thales CipherTrust Manager k170v v2.0.0 migration to client certificate authentication key 230
Enabling encryption on the storage system 231
Migrating from global key to Client Certificate Authentication 234
CipherTrust Manager cluster configuration 236
Thales SafeNet KeySecure 8.11.1 migration to client certificate authentication key 236
Enabling encryption on the storage system 237
Migrating from global key to Client Certificate Authentication 240
Thales SafeNet KeySecure 8.11.1 client certificate authentication for key management 242
Enabling encryption on the storage system 242

Reference 246
Command-line reference (12.3.2.x) 246
Overview of the command-line interface (CLI) 247
Overview of the XCLI utility 247
Installing and starting the XCLI 248
Starting the XCLI on a Windows system 248
Starting the XCLI on a Linux system 249
Exiting an interactive XCLI session 249
Using the CLI 249
Interactive and basic modes 250
Interactive mode features 250
Understanding the command syntax 251
Command syntax quick reference 251
Syntax example 251
Using identification parameters 252

Setting user and password parameters 252
Identifying and configuring a storage system 253
Configuration parameters 253
Creating a configuration file 253
Certificate management 254
Display options 255
Using display options 255
Table display options 255
Running commands in batch mode 256
Displaying CLI help and the XCLI utility version 257
Host and cluster management commands 257
Adding a host to a cluster 258
Creating a cluster 259
Deleting clusters 260
Listing clusters 260
Removing a host from a cluster 261
Renaming clusters 262
Adding a port to a host 262
Defining a new host 263
Deleting a host 264
Listing hosts 265
Listing ports 266
Removing a port from a host 266
Renaming a host 267
Updating a host definition 268
Mapping a volume to a host or cluster 269
Listing the mapping of volumes to hosts, clusters, or domains 270
Setting the special type of hosts or clusters 271
Listing hosts/clusters to which a volume is mapped 272
Unmapping a volume from a host or cluster 273
Set the default idle time before unmapping a volume 274
Retrieving the default idle time before unmapping a volume 275
Creating a performance class 275
Deleting a performance class 275
Renaming a performance class 276
Listing details of performance classes 277
Adding a host to a performance class 277
Removing a host from its performance class 278
Adding a pool to a performance class 278
Removing a pool from its performance class 279
Adding a volume to a performance class 280
Removing a volume from its performance class 280
Adding a domain to a performance class 281
Removing a domain from its performance class 282
Setting the rate for a performance class 282
Listing host profiles 283
Updating the host profile 283
Removing the profile of the specified host 284
Enabling the host profiler 284
Disabling the host profiler 285
Volume management commands 285
Clearing reservations of a volume 286
Listing reservation keys 286
Listing volume reservations 287
Copying volumes 287
Creating a volume 289
Deleting a volume 291
Formatting a volume 292
Listing volumes 293
Listing a volume's extended attributes 295
Locking a volume 295
Renaming a volume 297
Resizing a volume 298
Unlocking a volume 300
Volume snapshot management commands 301
Changing a snapshot deletion priority 302
Creating a snapshot 303
Deleting a snapshot 305
Duplicating a snapshot 306

Formatting a snapshot 307
Listing snapshot information 308
Restoring a volume from a snapshot 309
Consistency group management commands 310
Adding a volume to a consistency group 311
Creating consistency groups 314
Deleting a consistency group 315
Listing consistency groups 316
Removing a volume from a consistency group 316
Renaming a consistency group 318
Creating a cross-system consistency group 319
Associating an existing consistency group with a cross-system consistency group definition 319
Removing an existing consistency group from a cross-system consistency group definition 320
Adding a remote system name to a cross-system consistency group definition 321
Removing a remote system name from a cross-system consistency group definition 322
Listing cross-system consistency group definitions 322
Retrieving remote systems in a specified cross-system consistency group 323
Deleting a cross-system consistency group 324
Listing cross-system consistency group definitions 324
Snapshot set management commands 325
Snapshotting a consistency group 325
Changing a snapshot group deletion priority 327
Deleting a snapshot group 328
Disbanding a snapshot group 329
Duplicating a snapshot group 329
Formatting a snapshot group 330
Listing snapshot groups 331
Locking a snapshot group 332
Renaming a snapshot group 333
Restoring a consistency group from a snapshot group 333
Unlocking a snapshot group 334
Setting a snapshot group descriptor 335
Returning a snapshot group's descriptor 336
Storage pool management commands 336
Moving a consistency group between storage pools 337
Changing the pool limitation, performance class, or threshold parameters 338
Changing pool settings for snapshots 339
Creating storage pools 340
Deleting a storage pool 342
Listing storage pools 342
Renaming a storage pool 343
Resizing a storage pool 344
Moving a volume between storage pools 345
System management commands 347
Displaying the values of configuration parameters 348
Setting configuration parameters 349
Testing the DNS 350
Displaying help 350
Printing the current maintenance urgency 351
Shutting down the system 352
Listing the operational state 352
Displaying system usage and data reduction statistics 353
Displaying information about physical and effective capacity 354
Displaying information about effective capacity 355
Displaying system capacity thresholds 355
Changing a system capacity threshold 356
Resuming the system's normal operation 357
Displaying the current time 357
Setting the system's time 358
Listing optional time zones 358
Setting the time zone 359
Printing the current system version 359
Displaying the values of VPD parameters 360
Setting VPD parameters 360
Displaying the system's MIB file 362
Retrieving the electronic license acceptance status 363
Retrieving a fragment of the electronic license file 363
Accepting the electronic license agreement 364
Enabling command auditing 364

Disabling command auditing 365
Displaying the command audit state 365
Configuring audit servers 366
Checking the command audit state 366
Retrieving the list of Flash VDisks 367
Enabling CIM service 367
Disabling the CIM service 368
Displaying the CIM service state 368
Remote target connectivity commands 369
Setting the threshold of a link disruption duration that triggers an event 369
Updating the target's mirroring configuration 370
Activating connectivity to a remote target 370
Deactivating connectivity to a remote target 371
Defining connectivity to a remote target 372
Deleting connectivity to a remote target 373
Listing target connectivity definitions 374
Defining a remote target 375
Deleting a remote target 376
Listing remote targets 376
Allowing remote mirroring access 377
Activating a port 378
Adding a new port to a remote target 378
Deactivating a port 379
Deleting a port from a remote system 380
Listing the ports of a remote target 381
Renaming a remote target 381
Updating the target configuration 382
Adding a Quorum Witness to a target 382
Removing a Quorum Witness from a target 383
Remote mirroring commands 383
Canceling a snapshot mirror (ad hoc sync job) 384
Creating a snapshot mirror (ad hoc sync job) 385
Activating mirroring 389
Changing the RPO for local or remote system 391
Changing the designation of mirroring peers 392
Changing a mirroring schedule for remote slave peers 394
Changing the role of a mirrored volume 395
Changing a mirroring schedule for local peers 397
Creating a mirroring definition 398
Deactivating mirroring 403
Deleting a remote mirroring definition 404
Viewing the mirroring status 406
Obtaining statistics on past sync jobs 408
Switching roles between master and slave 409
Retrieving RPO thresholds 411
Setting an RPO threshold 411
Changing the interval of a schedule 412
Creating a schedule object 413
Triggering a schedule 414
Deleting a schedule object 415
Listing a schedule object 416
Renaming a schedule 416
Viewing sync job status 417
HyperSwap commands 417
Creating a HyperSwap relation 418
Viewing the status of HyperSwap volumes and consistency groups 423
Activating a HyperSwap relation 424
Deactivating a HyperSwap relation 427
Deleting a HyperSwap relation 428
Switching roles between Master and Slave volumes 429
Changing a peer role in a HyperSwap volume 431
Restoring the availability of a Master volume 433
Creating a HyperSwap volume snapshot (ad hoc sync job) 434
Changing the designation of HyperSwap relation peers 436
Enabling high availability of a HyperSwap relation 438
Disabling high availability of a HyperSwap relation 439
Converting a HyperSwap relation into a sync mirror 439
Converting a sync mirror into a HyperSwap relation 441
Creating a new Quorum Witness 443

Listing Quorum Witnesses 444
Updating a Quorum Witness definition 445
Renaming a Quorum Witness 446
Deleting a Quorum Witness 446
Activating a Quorum Witness 447
Deactivating a Quorum Witness 448
Listing the Quorum Witness connection status 449
Getting Quorum Witness information 450
Data migration commands 451
Activating data migration 451
Deactivating data migration 452
Defining data migration configuration 452
Deleting a data migration process 454
Listing data migration statuses 455
Testing the data migration definition 456
IBM Hyper-Scale Mobility commands 456
Creating an IBM Hyper-Scale Mobility relation 457
Activating a volume migration 460
Deactivating IBM Hyper-Scale Mobility migration 461
Aborting a defined or activated IBM Hyper-Scale Mobility process 462
Moving the IBM Hyper-Scale Mobility source volume to a Proxy state 463
Deleting an IBM Hyper-Scale Mobility relation 464
Listing the IBM Hyper-Scale Mobility status 465
Event handling commands 466
Generating a custom event 467
Generating a CSS product event 467
Defining a new event notification destination 468
Deleting a destination 470
Listing event notification destinations 471
Renaming a destination 472
Testing a destination 472
Updating an event notification destination 473
Adding a destination to a destination group 475
Creating a destination group 476
Updating an event notification destination group 477
Deleting a destination group 478
Listing destination groups 478
Removing a destination from a destination group 479
Renaming a destination group 479
Clearing alerting events 480
Listing events 481
Listing uncleared alerting events 482
Setting the threshold for event notification 483
Listing thresholds 483
Activating a rule 484
Creating event notification rules 485
Deactivating a rule 487
Deleting event notification rules 487
Listing event notification rules 488
Renaming event notification rules 489
Updating an event notification rule 489
Defining an SMS gateway 491
Deleting an SMS gateway 492
Listing SMS gateways 493
Prioritizing SMS gateways 493
Renaming an SMS gateway 494
Updating an SMS gateway 495
Defining a new SMTP gateway 496
Deleting an SMTP gateway 496
Listing SMTP gateways 497
Prioritizing SMTP gateways 498
Renaming an SMTP gateway 498
Updating an SMTP gateway 499
Generating an XMPNS admin control event 500
IP configuration commands 500
Updating the configuration of the Ethernet port 501
Showing the status and configuration of Ethernet ports (deprecated) 502
Showing the status and configuration of Ethernet ports 503
Displaying Ethernet (physical) port statistics in the system 505

Creating a new IP interface 506
Deleting IP interfaces 507
Listing IP interface configuration 508
Listing IP interface addresses 509
Displaying IP interface statistics in the system 509
Renaming an IP interface 510
Printing the ARP database of an IP interface 510
Testing the traceroute to a remote IP 511
Testing the traceroute to a remote IPv6 address 512
Updating an IP interface 512
Updating VLAN priority code points 514
Defining a new IPSec connection 514
Updating an existing IPSec connection 515
Removing an existing IPSec connection 516
Listing IPSec connections 517
Listing IPSec tunnels 517
Connecting to a support center 518
Defining a support center 519
Deleting a support center 519
Disconnecting from a support center 520
Listing support centers 520
Listing statuses of all support centers 521
Configuring the support center connection to enable automatic connect on restart 521
Listing the configuration of support center automatically connect 522
Creating a new IP access group 523
Removing an address from an IP access group 523
Adding a new address to an IP access group 524
Deleting an existing IP access group 524
Renaming an existing IP access group 525
Listing IP access groups 525
Listing IP access group addresses 526
PKI configuration commands 527
Listing PKI items 527
Generating a certificate signing request 527
Generating a private key and CSR 528
Deleting a PKI content 529
Changing a PKI symbolic name 529
Importing a signed certificate 530
Importing a PKCS#12 certificate 531
Displaying the details of a signed certificate 532
Updating a PKI certificate or services 533
InfiniBand commands 534
Listing the configured InfiniBand ports 534
Listing data counters for the enabled InfiniBand switch ports 535
Listing error counters for enabled InfiniBand switch ports 535
Listing the status of the enabled InfiniBand switch ports 536
Listing data counters for the enabled InfiniBand HCA ports 537
Listing error counters for the enabled InfiniBand HCA ports. 538
Listing the statuses of the enabled InfiniBand HCA ports 538
Listing the configured InfiniBand switches 539
Listing the configured InfiniBand switch management addresses 541
Listing the configured InfiniBand switch firmware versions 541
Listing the configured InfiniBand switch power values 542
Listing the configured InfiniBand switch voltage values 543
Listing the configured InfiniBand switch temperature values 544
Listing the configured InfiniBand switch fan parts 545
Listing the configured InfiniBand switch PSUs 545
Listing the configured InfiniBand switch BBUs 546
Listing the configured InfiniBand switch fans 547
Access control commands 548
Adding an access control definition 550
Deleting an access control definition 550
Listing access control definitions 551
Adding an LDAP server definition 552
Testing an LDAP configuration 553
Listing LDAP configuration parameters 554
Configuring LDAP in the system 555
Listing LDAP servers defined in the system 557
Listing LDAP server users 558

Listing LDAP-based authentication mode 558
Enabling or disabling LDAP-based authentication mode 559
Updating an LDAP server definition 560
Removing an LDAP server definition 561
Launching the ldapsearch utility 561
Defining a new user 563
Deleting a user 564
Adding users to user groups 565
Creating user groups 565
Deleting a user group 566
Listing user groups 567
Removing a user from a user group 568
Renaming user groups 568
Updating a user group 569
Listing users 570
Renaming users 571
Updating a user definition 571
Creating a new domain 573
Updating a domain definition 574
Renaming a domain 575
Deleting a domain 576
Listing domains 577
Listing users per domain 578
Listing objects in domains 579
Listing the global domain 579
Attaching an object to a domain 580
Disassociating object from a domain 581
Associating users to a domain 583
Removing a user from a domain 584
Adding a pool to a domain 584
Removing a pool from a domain 585
Moving a pool from one domain to another 586
Setting the domain attribute 587
Setting domain related policies 588
Displaying domain-related policies 589
Specifying a user associated with IBM Hyper-Scale Manager 589
Retrieving the user associated with the IBM Hyper-Scale Manager 590
Setting the application administrator's scope of commands 590
Getting the application administrator's scope of commands 591
Fibre Channel and iSCSI configuration and status commands 591
Discovering FC hosts 592
Changing FC port configuration 592
Listing FC ports 593
Listing FC Port Tests 595
Starting FC Port Test 596
Aborting FC Port Test 597
Listing connectivity to hosts 597
Flash enclosure maintenance commands 598
Listing Flash enclosure status 599
Listing a Flash canister status 600
Listing a Flash card status 601
Retrieving the list of Flash fans 602
Retrieving the list of Flash IB adapters 602
Retrieving the Flash control connectivity list 603
Retrieving the list of Flash PSUs 604
Retrieving the list of Flash PIBs 605
Retrieving the list of Flash LED cards 606
Listing Flash BBU components 606
Retrieving the serial number of a flash enclosure 1S 607
Suspending the automatic Flash BBU calibration 608
Hardware maintenance commands 608
Monitoring the rebuilding or redistribution processes 609
Displaying the system's average power consumption 609
Getting the values for calculating the system's average power consumption 610
Displaying the system's average temperature 610
Enabling XIV Support access 611
Disabling XIV Support access 612
Displaying the XIV Support window 612
Listing system components 613

Listing module configuration 614
Listing the internal temperature of modules 616
Listing boot media devices in the system 617
Listing vault devices in the system 619
Listing BBUs in the system 620
Listing PSUs in the system 622
Listing compression adapters in the system 623
Listing fans in a module 624
Listing NICs in the system 625
Listing DIMMs in the modules 626
Listing CPUs in the modules 629
Listing InfiniBand HCA Adapters in the storage system 630
Listing CNA adapters in the system 631
Listing module LEDs in the system 632
Listing data disk devices in the system 633
Listing service statuses 635
Listing system components that require service 635
Listing trace snapshot on a module. 636
Creating a trace snapshot 637
Statistics commands 637
Getting performance statistics 637
Retrieving usage history 640
Metadata commands 641
Setting metadata 641
Deleting metadata 642
Listing metadata 643
Setting user metadata 644
Listing user metadata 644
Deleting user metadata 645
Encryption enablement and support commands 646
Disabling encryption 646
Enabling encryption 648
Defining a key server 649
Removing a key server 650
Displaying key server status 651
Checking key server status 652
Obtaining a new master key 653
Renaming a key server 654
Changing key server properties 654
Entering a recovery key 656
Generating recovery keys 657
Retrieving the security administrator's recovery key 658
Rekeying the security administrators 659
Displaying recovery key status 660
Recovering key verification 661
Recovering key share information 662
Finishing the recovery process 663
Obtaining a new master key 663
Changing the key management scheme 665
Viewing the key scheme 666
Multi-site high availability and disaster recovery commands 666
Viewing the Multi-site relation status 666
Creating a Multi-site relation 668
Activating an asynchronous mirror within a Multi-site relation 670
Deleting a Multi-site relation 672
Changing roles in a Multi-site relation 673
Adding a Standby mirror to a Multi-site relation 675
Switching roles between Master and SMaster 676
Security configuration commands 678
Listing configuration parameters for a communication protocol 678
Setting configuration parameters for a communication protocol 679
Events 679
Return Codes 749
REST API reference 749
Open API reference 749
Introduction 749
Storage Management Initiative Specifications (SMI-S) overview 750
Common Information Model (CIM) agent overview 750
Common Information Model (CIM) concepts 751
Common Information Model (CIM) agent components 752

Open API element definitions 752
Common Information Model (CIM) agent security 753
Installation and configuration 753
Multitenancy feature support information 754
Communication concepts and methods 754
Common Information Model (CIM) agent communication concepts 754
Common Information Model (CIM) agent communication methods 754
Obtaining a single class from the target namespace 756
Obtaining a single instance from the target namespace 756
Deleting a single instance from the target namespace 756
Creating an instance in the target namespace 757
Modifying an existing instance in the target namespace 757
Enumerating classes within a defined target namespace 757
Enumerating the names of subclasses of a class defined within the target namespace 758
Enumerating instances of a defined class within the target namespace 758
Enumerating the names of instances of a class within a target namespace 759
Processing a query against the target namespace 759
Enumerating classes or instances that are associated with a specific Common Information Model (CIM) object 759
Enumerating the names of classes or instances associated with a specific Common Information Model (CIM) object 760
Enumerating the association objects that refer to a specific target class or instance 760
Enumerating the names of the association objects that refer to a specific target class or instance 761
Retrieving a single property value from an instance in the target namespace 761
Setting a single property value within an instance in the target namespace 762
Retrieving a single qualifier declaration from the target namespace 762
Creating or modifying a qualifier declaration in the target namespace 762
Enumerating qualifier declarations from the target namespace 763
Common Information Model (CIM) agent communication methods that cannot be used 763
DeleteClass method (not supported) 763
CreateClass method (not supported) 763
ModifyClass method (not supported) 763
DeleteQualifier method (not supported) 764
Return error codes 764
Functional profiles, diagrams, and methods 765
Block Server performance profile 765
Block Server Performance object model 766
Block Server Performance methods 768
Obtaining performance statistics data 768
Obtaining volume or host statistics 768
Block Services profile 769
Block Services object model 769
Block Services methods 772
Creating or modifying a virtual pool or snapshot pool 773
Deleting a storage pool 773
Creating, modifying, and/or moving multiple volumes in a single method call 773
Creating, modifying, and/or moving a single volume 774
Deleting a single volume 774
Deleting multiple volumes concurrently 774
Determining sizes to use to create elements from the primordial storage pool 775
Determining sizes to use to create elements from the virtual storage pool 775
Creating a concrete storage pool 775
Creating a storage volume 775
iSCI Target Ports profile 775
iSCI Target Ports object model 776
iSCSI Target Ports methods 777
Creating an IP endpoint 777
Parameters 778
Modifying an IP endpoint 778
Parameters 778
Deleting an IP endpoint 779
Parameters 779
Masking and Mapping profile 779
Masking and Mapping object model 781
Masking and Mapping methods 782
Creating a storage hardware ID 782
Deleting a storage hardware ID 783
Creating a hardware ID collection 783
Adding hardware IDs to a collection 783
Assigning volumes to a storage hardware ID 783
Removing access to volumes from a storage hardware ID 784

Deleting an existing protocol controller 784
Example configuration procedures using the Masking and Mapping profile 784
Indication profile 785
Replication Services profile 786
Replication Services object model 786
Replication Services methods 789
Group management 789
Creating a consistency group 790
Deleting a consistency group 790
Adding volumes to a consistency group 790
Removing volumes from a consistency group 791
Replication management 791
Creating a snapshot or clone 792
Creating a snapshot group for a consistency group 793
Restoring a snapshot to a volume or manipulating a mirror consistency group 793
Modifying a list of synchronizations using a batch operation 794
Retrieving target elements 794
Retrieving snapshot or mirror relationships 794
Retrieving strings of references to snapshots, snapshot groups, or mirror relationships 794
Retrieving peer systems 794
Creating new storage objects that are replicas of specified source storage objects 795
Replication Services indications 795
Job Control profile 795
Job Control object model 796
Thin provisioning profile 796
Thin Provisioning object model 797
Thin Provisioning methods 797
Creating an IBMTSDS_VirtualPoolSetting instance 797
Modifying the properties of a newly created IBMTSDS_VirtualPoolSetting instance 798
Deleting a new IBMTSDS_VirtualPoolSetting instance 798
Creating or modifying a thin provisioning virtual pool 798
Creating a thin provisioning pool 799
Thin Provisioning indications 799
Conformance tests 799

Remote support 799
Remote support and proactive support 800
Remote Support Proxy 800
Introducing the IBM Remote Support Proxy utility 800
Compatibility and requirements 800
Installation 801
Installing and configuring for the first time 801
Installing the software package 801
Configuring the Remote Support Proxy service 802
Starting the Remote Support Proxy service 804
Configuring storage systems 804
Upgrading and configuring an existing installation 805
Uninstalling the Remote Support Proxy 805
Management 806
Starting and stopping the Remote Support Proxy service 806
Monitoring status 806
Updating remote-support server information 806

Publications and related information 807

Glossary 808
A 808
B 808
C 809
D 810
E 811
F 812
G 813
H 813
I 814
J 815
L 815
M 816
N 816
O 817
P 817
Q 818

R 818
S 819
T 821
U 821
V 821
W 822
Z 822

IBM FlashSystem A9000R welcome page

This page provides an entry point to product information about IBM FlashSystem A9000R. Select a link in the dropdown menu above to find documentation for a specific version of this product. Generated links to developerWorks, Redbooks, and white papers are also provided when articles about the product are available. On this page, you can find entry points to sites that help you learn more about this product and other IBM products. There are also links to web sites that help you find support and stay current.

Cross-system feature use in the IBM Spectrum Accelerate family

The tables in the following sections specify which functional features are available when used across different IBM Spectrum Accelerate family storage systems and microcode versions.

The IBM Spectrum Accelerate family includes:

- IBM FlashSystem® A9000 (12.x.x) (FlashSystem A-line)
- IBM FlashSystem A9000R (12.x.x) (FlashSystem A-line)
- IBM® XIV® Storage System Gen3 (11.x.x) and IBM XIV Storage System (10.2.4.x)
- IBM Spectrum Accelerate deployable software (11.5.x) (deployed on 3–15 VMware ESXi hosts)

Important: For full IBM Spectrum Accelerate end of support information, see https://ibm.biz/BdfWcQ and https://ibm.biz/BdfWcP.

Note: The listed functional features are supported by using:

- IBM Hyper-Scale Manager management GUI 5.5.0 or later
- IBM Extended Command-Line Interface (XCLI) utility 5.2.0 or later

- HyperSwap
- Synchronous mirroring (two systems)
- Asynchronous mirroring (two systems)
- Multi-site HA/DR (High Availability & Disaster Recovery) for system software 12.x.x
- Multi-site mirroring ("3-way") for XIV Gen3 and Spectrum Accelerate deployable software
- IBM Hyper-Scale Mobility
- Data migration

HyperSwap

HyperSwap® delivers highly-available, non-disruptive storage service, through partial or complete system failures and disasters, in the same data center and between metro-distant data centers.

HyperSwap high availability is based on active-active pairing of storage systems per volume or per consistency group. Each volume or consistency group pair uses synchronous replication to keep both systems updated at all times.

Important: Use of HyperSwap requires Fibre Channel (FC) connectivity between the systems.
For more information, see High availability with HyperSwap.

From Primary/To Secondary	To IBM XIV Gen2	To IBM XIV Gen3	To IBM Spectrum Accelerate	To IBM FlashSystem A9000 12.1.0 or later	To IBM FlashSystem A9000R 12.1.0 or later
From IBM XIV Gen2					
From IBM XIV Gen3					
From IBM Spectrum Accelerate					
From IBM FlashSystem A9000 12.1.0 or later				✓	✓
From IBM FlashSystem A9000R 12.1.0 or later				✓	✓

Synchronous mirroring (two systems)

The process of ensuring that both storage systems contain identical data at all times is called *remote mirroring*. Synchronous remote mirroring provides continuous availability of critical information in the case of a disaster scenario. For more information, see Synchronous remote mirroring.

From Primary/To Secondary	To IBM XIV Gen2	To IBM XIV Gen3	To IBM Spectrum Accelerate	To IBM FlashSystem A9000 12.0.1 or later	To IBM FlashSystem A9000R 12.0.1 or later
From IBM XIV Gen2	✓	✓			
From IBM XIV Gen3	✓	✓	✓		

From Primary/To Secondary	To IBM XIV Gen2	To IBM XIV Gen3	To IBM Spectrum Accelerate	To IBM FlashSystem A9000 12.0.1 or later	To IBM FlashSystem A9000R 12.0.1 or later
From IBM Spectrum Accelerate		✓	✓		
From IBM FlashSystem A9000 12.0.1 or later				✓	✓
From IBM FlashSystem A9000R 12.0.1 or later				✓	✓

Asynchronous mirroring (two systems)

Asynchronous mirroring enables you to attain high availability of critical data through a process that asynchronously replicates updated data from a primary storage peer to a remote, secondary peer. For more information, see Asynchronous remote mirroring.

From Primary/To Secondary	To IBM XIV Gen2	To IBM XIV Gen3 11.0.0 or later	To IBM XIV Gen3 11.6.2.a or later	To IBM Spectrum Accelerate	To IBM FlashSystem A9000 12.0.1 or later	To IBM FlashSystem A9000R 12.0.1 or later	To IBM FlashSystem A9000 12.2.0 or later	To IBM FlashSystem A9000R 12.2.0 or later
From IBM XIV Gen2	✓	✓	✓					
From IBM XIV Gen3 11.0.0 or later	✓	✓	✓	✓				
From IBM XIV Gen3 11.6.2.a or later	✓	✓	✓	✓			✓	✓
From IBM Spectrum Accelerate		✓	✓	✓				
From IBM FlashSystem A9000 12.0.1 or later					✓	✓	✓	✓
From IBM FlashSystem A9000R 12.0.1 or later					✓	✓	✓	✓
From IBM FlashSystem A9000 12.2.0 or later			✓		✓	✓	✓	✓
From IBM FlashSystem A9000R 12.2.0 or later			✓		✓	✓	✓	✓

Multi-site HA/DR (High Availability & Disaster Recovery) for system software 12.x.x

The Multi-site HA/DR function allows customers to deploy high availability and disaster recovery solutions over multiple sites, keeping three concurrent copies of data.

Important: Use of Multi-site HA/DR requires that the HyperSwap relation have Fibre Channel (FC) connectivity between the systems, while the Asynchronous relations (active and standby) can be either over FC or over iSCSI.
For more information, see Multi-site mirroring.

From Primary/To Secondary	To IBM FlashSystem A9000 12.3.0 or later	To IBM FlashSystem A9000R 12.3.0 or later
From IBM FlashSystem A9000 12.3.0 or later	✓	✓
From IBM FlashSystem A9000R 12.3.0 or later	✓	✓

Multi-site mirroring ("3-way") for system software 11.x.x

Multi-site mirroring (also known as "3-way mirroring") is an IBM XIV Storage System technology that allows customers to set disaster recovery solutions over multiple sites, keeping three copies of their data. For more information, see Multi-site mirroring.

From Primary/To Secondary	To IBM XIV Gen3	To IBM Spectrum Accelerate

From Primary/To Secondary	To IBM XIV Gen3	To IBM Spectrum Accelerate
From IBM XIV Gen3 11.6.0 or later	✓	
From IBM Spectrum Accelerate		

IBM Hyper-Scale Mobility

IBM Hyper-Scale Mobility enables non-disruptive migration of volumes from one storage system to another. IBM Hyper-Scale Mobility helps achieve storage management objectives that are otherwise difficult to address, like:

- Migrating data out of an over-provisioned system.
- Migrating all the data from a system that will be decommissioned or re-purposed.
- Migrating data to another storage system to achieve adequate (lower or higher) performance, or to load-balance systems to ensure uniform performance.
- Migrating data to another storage system to load-balance capacity utilization.

Important: Use of IBM Hyper-Scale Mobility requires Fibre Channel (FC) connectivity between the systems, except for on IBM Spectrum Accelerate storage systems which require iSCSI connectivity.
For more information, see IBM Hyper-Scale Mobility.

From Primary/To Secondary	To IBM XIV Gen3	To IBM Spectrum Accelerate 11.5.3 or later	To IBM FlashSystem A9000 12.0.1 or later	To IBM FlashSystem A9000R 12.0.1 or later	To IBM FlashSystem A9000 12.2.0 or later	To IBM FlashSystem A9000R 12.2.0 or later	To IBM FlashSystem A9000 12.2.1 or later	To IBM FlashSystem A9000R 12.2.1 or later
From IBM XIV Gen3	✓						✓	✓
From IBM Spectrum Accelerate 11.5.3 or later		✓						
From IBM FlashSystem A9000 12.0.1 or later			✓		✓		✓	
From IBM FlashSystem A9000R 12.0.1 or later								
From IBM FlashSystem A9000 12.2.0 or later			✓		✓	✓	✓	✓
From IBM FlashSystem A9000R 12.2.0 or later					✓	✓	✓	✓
From IBM FlashSystem A9000 12.2.1 or later			✓		✓	✓	✓	✓
From IBM FlashSystem A9000R 12.2.1 or later					✓	✓	✓	✓

Data migration

Note: Data migration can be done from any SCSI storage device to all storage systems of the IBM Spectrum Accelerate family.

From Primary/To Secondary	To IBM XIV Gen2	To IBM XIV Gen3	To IBM Spectrum Accelerate 11.5.3 or later	To IBM FlashSystem A9000 12.0.0 or later	To IBM FlashSystem A9000R 12.0.0 or later
From IBM XIV Gen2	✓	✓	✓	✓	✓
From IBM XIV Gen3	✓	✓	✓	✓	✓
From IBM Spectrum Accelerate 11.5.3 or later	✓	✓	✓	✓	✓
From IBM FlashSystem A9000 12.0.0 or later	✓	✓	✓	✓	✓

From Primary/To Secondary	To IBM XIV Gen2	To IBM XIV Gen3	To IBM Spectrum Accelerate 11.5.3 or later	To IBM FlashSystem A9000 12.0.0 or later	To IBM FlashSystem A9000R 12.0.0 or later
From IBM FlashSystem A9000R 12.0.0 or later	✓	✓	✓	✓	✓
From SCSI Storage Array	✓	✓	✓	✓	✓

Platform and application integration

IBM® storage systems are integrated with various independent software vendor (ISV) platforms and applications on the host side. To better incorporate IBM storage system with these ISV platforms and applications, IBM provides different software solutions for server virtualization, geo-clustering, failover, storage management, host connectivity, and more.

The following links lead to the online documentation and full version information of the currently available host-side solutions for the various IBM storage products. To know more about which system storage versions are supported, see the release notes for the individual platform or application.

- IBM DS8000® Family product integration
- IBM FlashSystem® A9000 and A9000R product integration
- IBM Spectrum® Virtualize family and Storwize® product integration
- IBM XIV® Storage System product integration
- IBM Spectrum Accelerate integration
- IBM Storwize V7000 Unified product integration

IBM DS8000 Family products integrate with the following

- IBM block storage CSI driver
- IBM Spectrum Connect
- IBM Storage Enabler for Containers
- IBM Storage Management Pack for Microsoft System Center Operations Manager (SCOM)
- IBM Storage Driver for OpenStack
- IBM Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service
- IBM DS8000 Family Storage Replication Adapter (SRA)

IBM FlashSystem A9000 and A9000R integrate with the following

- IBM block storage CSI driver
- IBM Spectrum Connect
- IBM Storage Enabler for Containers
- IBM Storage Management Pack for Microsoft System Center Operations Manager (SCOM)
- IBM Storage Driver for OpenStack
- IBM Spectrum Accelerate Family Provider for Microsoft Volume Shadow Copy Service documentation
- IBM Spectrum Accelerate Family Storage Replication Adapter (SRA)

In addition, for IBM Storage Host Attachment Kit (HAK) product information, see IBM Storage Host Attachment Kit. IBM Storage Host Attachment Kit is no longer supported but may still work on your operating system. See the product documentation for more information. For more information about direct host attachment connectivity support, see the latest Configuring section of your FlashSystem A9000 or A9000R documentation.

IBM Spectrum Virtualize family, including SAN Volume Controller (SVC) and FlashSystem® family members built with Spectrum® Virtualize (including FlashSystem 5xxx, 7xxx, 9xxx), and Storwize products, integrate with the following

- IBM block storage CSI driver
- IBM Spectrum Connect
- IBM Storage Enabler for Containers
- IBM Storage Management Pack for Microsoft System Center Operations Manager (SCOM)
- IBM Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service
- IBM Spectrum Virtualize Family Storage Replication Adapter (SRA)
- ODF FlashSystem driver

In addition, the IBM Spectrum Virtualize family also integrates with the following plugins:

- Spectrum Virtualize family volume driver for Cinder supports various OpenStack versions of Storwize family and IBM FlashSystem systems. For more information, see the Reference >OpenStack support summary page of your specific storage product. For release notes and user information, see OpenStack Documentation.
- Ansible® Collections for IBM Spectrum Virtualize. For the latest Ibm.Spectrum_Virtualize Ansible Collection documentation, see Ibm.Spectrum_Virtualize Collection on Ansible Documentation.

IBM XIV Storage System integrates with the following

- IBM Storage Enabler for Containers
- IBM Storage Enabler for Windows Failover Clustering
- IBM Spectrum Connect
- IBM Storage Driver for OpenStack
- IBM Storage Management Pack for Microsoft System Center Operations Manager (SCOM)
- IBM Spectrum Accelerate Family Provider for Microsoft Windows Volume Shadow Copy
- IBM Spectrum Accelerate Family Storage Replication Adapter (SRA)

In addition, for IBM Storage Host Attachment Kit (HAK) product information, see IBM Storage Host Attachment Kit. IBM Storage Host Attachment Kit is no longer supported but may still work on your operating system. See the product documentation for more information.

IBM Spectrum Accelerate integrates with the following

Important: For full IBM Spectrum Accelerate end of support information, see https://ibm.biz/BdfWcQ and https://ibm.biz/BdfWcP.

- IBM Storage Enabler for Containers
- IBM Spectrum Connect
- IBM Storage Driver for OpenStack
- IBM Storage Management Pack for Microsoft System Center Operations Manager (SCOM)
- IBM Spectrum Accelerate Family Provider for Microsoft Windows Volume Shadow Copy
- IBM Spectrum Accelerate Family Storage Replication Adapter (SRA)

In addition, for IBM Storage Host Attachment Kit (HAK) product information, see IBM Storage Host Attachment Kit. IBM Storage Host Attachment Kit is no longer supported but may still work on your operating system. See the product documentation for more information.

IBM Storwize V7000 Unified integrates with the following

- IBM Spectrum Virtualize Family Storage Replication Adapter (SRA)

GDPR and IBM systems

Refer to the following white paper for details about the impact of the General Data Protection Regulation (GDPR) on systems and storage.

Impact of GDPR on systems and storage

To view a PDF file, you need the Adobe Acrobat Reader. You can download it at no charge from the Adobe web site.

What's new

The following table lists the dates and nature of updates to the published information for IBM FlashSystem® A9000R 12.3.x.

Date	Nature of updates to the published information

18 January 2023	• New release for IBM FlashSystem A9000R 12.3.2.e.
	  ◦ First edition release notes for IBM FlashSystem A9000R 12.3.2.e
	  ◦ First edition of the CLI Reference Guide for IBM FlashSystem A9000R 12.3.2.e.
	  ◦ For information about what's new in this version of IBM FlashSystem A9000R 12.3.2.e, see What's new in 12.3.2.e.
	• Second edition release notes for IBM Remote Support Proxy utility 1.3.2.2, adding new known issue PXY-235418.
	  ◦ For information about what's new in this version, see What's new in 1.3.2.2.
	• For the PDF version of the release notes, see the release notes PDF publications page.
	• For the PDF version of the other relevant publications, see the PDF publications page.

15 September 2022	• Updated the Configuring key servers section.

2 August 2022	• Updated the Configuring key servers section. The existing sections have been updated and the following key server information has been added:
	  ◦ Thales CipherTrust Manager k170v v2.3.x and 2.6.x migration to client certificate authentication key
	  ◦ Thales CipherTrust Manager k170v v2.2.x migration to client certificate authentication key
	• Second edition Product Overview, adding CipherTrust Manager k170v 2.6 support. For more information, see Encryption key management schemes.
	• For the PDF version of this and other relevant publications, see the PDF publications page.

13 July 2022	• Updated the Configuring host attachment section. This update adds the following host attachment information:
	  ◦ Preparing for host attachment
	  ◦ Defining the host on the storage system
	  ◦ Monitoring host connectivity
	  ◦ Configuring Linux host attachment
	  ◦ Configuring Windows host attachment
	  ◦ Configuring Solaris host attachment

5 July 2022	• Ninth edition user guide for IBM Remote Support Proxy utility 1.3.x. This edition updates the SERVER_LIST and ServerAddress parameter information. For more information, see Configuring the Remote Support Proxy service • For the PDF version of this and other relevant publications, see the PDF publications page.
16 June 2022	• Updated Configuring key servers section, adding information for Thales SafeNet KeySecure 8.11.1.
9 June 2022	• First edition release notes for XCLI Utility 5.5.4.1. This edition fixes issue PUI-17822. For more information, see the Change log. • For the PDF version of the release notes, see the release notes PDF publications page.

Date	Nature of updates to the published information
16 May 2022	• First edition release notes for XCLI Utility 5.5.4. This edition replaces 5.2.0.a. • For information about what's new in this version, see What's new in 5.5.4. • New edition CLI Reference guide with minor updates to XCLI Utility installation information. • For the PDF version of the release notes, see the release notes PDF publications page.
12 May 2022	• First edition release notes for IBM Remote Support Proxy utility 1.3.2.2. • For information about what's new in this version, see What's new in 1.3.2.2. • For the PDF version of the release notes, see the release notes PDF publications page.
29 March 2022	• First edition release notes for IBM FlashSystem A9000R 12.3.2.d. • New edition of the CLI Reference Guide for IBM FlashSystem A9000R 12.3.2.d. • New edition Deployment Guide, updating Host network connectivity and multipathing configurations. • New edition Product Overview, updating Host system attachment. • Added new Configuring section for additional host attachment and key server configuration. • For information about what's new in this version of IBM FlashSystem A9000R 12.3.2.d, see What's new in 12.3.2.d.
24 March 2022	• Second edition release notes for IBM FlashSystem A9000R 12.3.2.c. This edition changes the severity of SYS-316012 from low to HIPER. • For the PDF version of the release notes, see the release notes PDF publications page.
24 October 2021	• New edition of the IBM Spectrum® Accelerate Family HyperSwap® Quorum Witness 1.0.3 user guide. This edition updates Replacing a self-signed certificate on the Quorum Witness.

18 October 2021	• New edition of the Product Overview. This edition includes updates to the following sections:
	  ◦ High availability with HyperSwap
	  ◦ Mirroring resources

17 August 2021	• Third edition release notes for IBM Spectrum Accelerate Family HyperSwap Quorum Witness 1.0.3, adding RHEL and CentOS 8.4 support. • For information about what's new in this version of IBM Spectrum Accelerate Family HyperSwap Quorum Witness 1.0.3, see What's new in 1.0.3. • For the PDF version of the release notes, see the release notes PDF publications page.
25 July 2021	• Seventh edition release notes for IBM Remote Support Proxy utility 1.3.2.1, adding PXY-235410 as a new known issue. • Second edition release notes for IBM Spectrum Accelerate Family HyperSwap Quorum Witness 1.0.3, adding QUOR-248 as a new known issue. • For the PDF version of the release notes, see the release notes PDF publications page.
2 February 2021	• New user guide and release notes for IBM Spectrum Accelerate Family HyperSwap Quorum Witness 1.0.3. • For information about what's new in this version of IBM Spectrum Accelerate Family HyperSwap Quorum Witness 1.0.3, see What's new in 1.0.3. • Fourth edition of the IBM Spectrum Accelerate Family HyperSwap Quorum Witness 1.0.2 release notes. This edition adds QUOR-268 as a new known issue. • For the PDF version of the release notes, see the release notes PDF publications page. • For the PDF version of the user guide, see the PDF publications page.

18 January 2021	• First edition release notes for IBM FlashSystem A9000R 12.3.2.c.
	• New edition of the CLI Reference Guide for IBM FlashSystem A9000R 12.3.2.c.
	• New release notes editions, adding SYS-316182 and SYS-316205 as new known issues:
	  ◦ Third edition release notes for IBM FlashSystem A9000R version 12.3.2.b
	  ◦ Second edition release notes for IBM FlashSystem A9000R version 12.3.2.a-3
	  ◦ Sixth edition release notes for IBM FlashSystem A9000R version 12.3.2.a-1
	• For the PDF version of the release notes, see the release notes PDF publications page.
	• For the PDF version of other relevant publications, see the PDF publications page.

27 December 2020	• Sixth edition release notes for IBM Remote Support Proxy utility version 1.3.2.1. This edition includes updates to the System requirements section. • For the PDF version of the release notes, see the release notes PDF publications page.
21 December 2020	• Second edition release notes for IBM FlashSystem A9000R version 12.3.2.b. • For the PDF version of the release notes, see the release notes PDF publications page.
5 November 2020	• Third edition of the IBM Spectrum Accelerate Family HyperSwap Quorum Witness 1.0.2 Release Notes. This edition updates the supported Linux versions and information for known issue QUOR-262. For the PDF version of the release notes, see the release notes PDF publications page.
2 September 2020	• Second edition of the IBM Spectrum Accelerate Family HyperSwap Quorum Witness 1.0.2 Release Notes. This edition adds QUOR-262 as a new known issue. • New edition of the IBM Spectrum Accelerate Family HyperSwap Quorum Witness 1.0.2 User Guide. This edition updates for restoring the Quorum Witness database. For the PDF version of this guide, see the PDF publications page.

Date	Nature of updates to the published information
27 July 2020	• New editions of the CLI Reference Guide for IBM FlashSystem A9000R 12.3.0-12.3.2 and 12.3.2.b, updating the descriptions of the following CLI commands description updates: **audit_enable**, **dest_define**, **protocol_config_list**, and **protocol_config_set**. • Eighth edition release notes for XCLI Utility 5.2.0. XCLI Utility no longer supports Red Hat® Enterprise Linux (RHEL) with 64-bit architecture. • For the PDF version of the release notes, see the release notes PDF publications page. • For the PDF version of other relevant publications, see the PDF publications page.
26 May 2020	• First edition release notes for IBM FlashSystem A9000R version 12.3.2.b. • New edition of the CLI Reference Guide for IBM FlashSystem A9000R version 12.3.2.b. • For the PDF version of the release notes, see the release notes PDF publications page. • For the PDF version of other relevant publications, see the PDF publications page.
13 May 2020	• First edition release notes for IBM FlashSystem A9000R version 12.3.2.a-3. • New edition of the Product Overview. This edition includes updates to the "Encryption key management schemes" section. • For the PDF version of the release notes, see the release notes PDF publications page. • For the PDF version of other relevant publications, see the PDF publications page.
17 March 2020	• New edition of the IBM Remote Support Proxy utility version 1.3.2.1 User Guide. This edition includes updates to the Configuring the Remote Support Proxy service section. For the PDF version of this guide, see the PDF publications page.
30 January 2020	• New edition of the IBM Spectrum Accelerate Family HyperSwap Quorum Witness version 1.0.2 User Guide. This edition includes minor updates. For the PDF version of this guide, see the PDF publications page. • Fifth edition release notes for IBM FlashSystem A9000R version 12.3.2.a-1, and sixth edition release notes for IBM FlashSystem A9000R version 12.3.2.a. These editions add information about the IBM block storage CSI driver support for FlashSystem A9000R. • For the PDF version of the release notes, see the release notes PDF publications page.
11 December 2019	• First edition of the IBM Spectrum Accelerate Family HyperSwap Quorum Witness version 1.0.2 User Guide. For the PDF version of this guide, see the PDF publications page.

27 November 2019	• New editions, added a new fixed issue, SYS-307987, to the 12.3.0 change log:
	  ◦ Fourth edition release notes for IBM FlashSystem A9000R versions 12.3.1 and 12.3.2.a-1
	  ◦ Fifth edition release notes for IBM FlashSystem A9000R versions 12.3.0, 12.3.2, and 12.3.2.a
	  ◦ Sixth edition release notes for IBM FlashSystem A9000R version 12.3.0.a
	• For the PDF version of the release notes, see the release notes PDF publications page.

26 September 2019	Third edition release notes for IBM FlashSystem A9000R versions 12.3.2 and 12.3.2.a, and second edition release notes for IBM FlashSystem A9000R version 12.3.2.a-1. These editions add a new fixed issue, SYS-310148, to the 12.3.2 change log. • For the PDF version of the release notes, see the release notes PDF publications page.
9 September 2019	• First edition release notes for IBM FlashSystem A9000R version 12.3.2.a-1. • Second edition release notes for IBM FlashSystem A9000R version 12.3.2.a. Added SYS-315146 to the 12.3.2.a change log. • New edition of the CLI Reference Guide for IBM FlashSystem A9000R version 12.3.2.a, updating the **protocol_config_list** and **protocol_config_set** commands. • For the PDF version of the release notes, see the release notes PDF publications page • For the PDF version of other relevant publications, see the PDF publications page.
19 August 2019	• First edition release notes for IBM FlashSystem A9000R version 12.3.2.a. • New edition of the CLI Reference Guide for IBM FlashSystem A9000R version 12.3.2.a. • First edition release notes for IBM Spectrum Accelerate Family HyperSwap Quorum Witness version 1.0.2. • For the PDF version of the release notes, see the release notes PDF publications page • For the PDF version of other relevant publications, see the PDF publications page.
31 July 2019	• Second edition release notes for IBM FlashSystem A9000R version 12.3.2. This edition includes IBM Spectrum Accelerate Family HyperSwap Quorum Witness installation requirement updates. • New edition of the Product Overview. This edition includes IBM Spectrum Accelerate Family HyperSwap Quorum Witness installation requirement updates. • New edition of the CLI Reference Guide. • Second edition release notes for IBM Spectrum Accelerate Family HyperSwap Quorum Witness version 1.0.1. This edition includes installation requirement updates. • New edition of the IBM Spectrum Accelerate Family HyperSwap Quorum Witness version 1.0.1 User Guide. This edition includes installation requirement updates. • Fifth edition release notes for IBM Remote Support Proxy utility version 1.3.2.1. This edition includes updates to the supported storage systems section. • New edition of the IBM Remote Support Proxy utility version 1.3.2.1 User Guide. This edition includes an update in the Configuring the Remote Support Proxy service section. • For the PDF version of the release notes, see the release notes PDF publications page. • For the PDF version of other relevant publications, see the PDF publications page.
15 May 2019	Release of the IBM FlashSystem A9000R user documentation set for version 12.3.2. • For the PDF version of the release notes, see the release notes PDF publications page • For the PDF version of other relevant publications, see the PDF publications page.

Date	Nature of updates to the published information
11 February 2019	Third edition release notes for IBM FlashSystem A9000R version 12.3.0 and fourth edition release notes for IBM FlashSystem A9000R version 12.3.0.a. These editions add a new fixed issue, SYS-308558, to the 12.3.0 change log. • For the PDF version of the release notes, see the release notes PDF publications page.
23 January 2019	Release of the IBM FlashSystem A9000R user documentation set for version 12.3.1. • For the PDF version of the release notes, see the release notes PDF publications page • For the PDF version of other relevant publications, see the PDF publications page.
15 January 2019	Third edition release notes for IBM XCLI Utility version 5.2.0.a and sixth edition release notes for IBM XCLI Utility version 5.2.0. In these editions, installation instructions for the XCLI utility were removed and other information was updated, such as supported microcode versions and XCLIU release version numbers. • For the PDF version of the release notes, see the release notes PDF publications page.
11 January 2019	Second edition release notes for IBM FlashSystem A9000R version 12.3.0 and third edition release notes for IBM FlashSystem A9000R version 12.3.0.a. These editions correct a typographical error in pool size (1 TB to 1 PB) in the removal of storage pool size limitation introduced by version 12.3.0. This information appears in the What's new section and the 12.3.0 change log of respective release notes. • For the PDF version of the release notes, see the release notes PDF publications page.
7 December 2018	Second edition release notes for IBM FlashSystem A9000R version 12.3.0.a. This edition updates the Affected version number for the fixed issue SYS-309775. • For the PDF version of the release notes, see the release notes PDF publications page
23 November 2018	First edition release notes for IBM FlashSystem A9000R version 12.3.0.a. • For the PDF version of the release notes, see the release notes PDF publications page
21 November 2018	Release of the IBM FlashSystem A9000R user documentation set for version 12.3.0. • For the PDF version of the release notes, see the release notes PDF publications page • For the PDF version of other relevant publications, see the PDF publications page.

Agreements and license information

These license agreements and program license information documents apply to IBM FlashSystem A9000R.

Table 1. Product agreements and license information

Title	Link
International Agreement for Acquisition of Software Maintenance (Not all software will offer Software Maintenance under this agreement.)	http://ibm.co/1fBmKPz
Other Internal Licensed Code	http://ibm.co/1kvABXE
International Program License Agreement and International License Agreement for Non-Warranted Programs	http://ibm.co/1pWmQ5U

Warranty information

The statement of limited warranty specifies the warranty period, type of warranty service, and service level.

Warranty Information

To view a PDF file, you need the Adobe Acrobat Reader. You can download it at no charge from the Adobe web site.

Notices

The legal notices pertain to IBM FlashSystem® A9000R documentation.

This information was developed for products and services offered in the US. This material may 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
USA

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 states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

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

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.

- **Trademarks**
- **Homologation statement**
  This product may not be certified in your country for connection by any means whatsoever to interfaces of public telecommunications networks. Further certification may be required by law prior to making any such connection. Contact an IBM representative or reseller for any questions.

Trademarks

IBM®, the IBM logo, and ibm.com® are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at Copyright and trademark information at www.ibm.com/legal/copytrade.shtml.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Microsoft and Windows 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.

Homologation statement

This product may not be certified in your country for connection by any means whatsoever to interfaces of public telecommunications networks. Further certification may be required by law prior to making any such connection. Contact an IBM representative or reseller for any questions.

Safety and environmental notices

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

- **Safety notices and labels**
  Review the safety notices and safety information labels before using this product.
- **Danger notices for IBM FlashSystem A9000R systems**
  Ensure that you understand the danger notices for IBM FlashSystem® A9000R systems.
- **Caution notices for IBM FlashSystem A9000R systems**
  Ensure that you understand the caution notices for IBM FlashSystem A9000R systems.
- **Special caution and safety notices**
  This information describes special safety notices that apply to FlashSystem A9000R. These notices are in addition to the standard safety notices supplied and address specific issues relevant to the equipment provided.
- **Environmental notices**
  This information contains all the required environmental notices for IBM® Systems products in English and other languages.

Safety notices and labels

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

IBM Systems safety notices and information

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

IBM Systems Safety Notices, G229-9054

The publication is organized into three sections:

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

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

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

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

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

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

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

Finding translated notices

Each safety notice contains an identification number. You can use this identification number to check the safety notice in each language. The list of notices that apply to this product are listed in the Special caution and safety notices and Environmental notices topics of this guide.

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

1. In the product documentation, look for the identification number at the end of each caution notice or each danger notice. In the following examples, the numbers (D002) and (C001) are the identification numbers.
   DANGER
   A danger notice indicates the presence of a hazard that has the potential of causing death or serious personal injury. (D002)
   CAUTION:
   A caution notice indicates the presence of a hazard that has the potential of causing moderate or minor personal injury. (C001)
2. Open the *IBM Systems Safety Notices*.
3. Under the language, find the matching identification number. Review the topics concerning the safety notices to ensure that you are in compliance.

To view a PDF file, you need Adobe Reader. You can download it at no charge from the Adobe website.

Danger notices for IBM FlashSystem A9000R systems

Ensure that you understand the danger notices for IBM FlashSystem A9000R systems.

Danger notices

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

DANGER: An electrical outlet that is not correctly wired could place hazardous voltage on the metal parts of the system or the devices that attach to the system. It is the responsibility of the customer to ensure that the outlet is correctly wired and grounded to prevent an electrical shock. (D004)
DANGER: When working on or around the system, observe the following precautions:
Electrical voltage and current from power, telephone, and communication cables are hazardous. To avoid a shock hazard: If IBM supplied the power cord(s), connect power to this unit only with the IBM provided power cord. Do not use the IBM provided power cord for any other product. Do not open or service any power supply assembly. Do not connect or disconnect any cables or perform installation, maintenance, or reconfiguration of this product during an electrical storm.

- 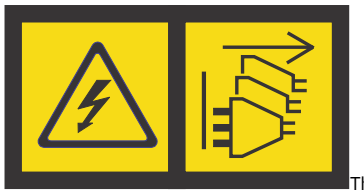
  The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords. For AC power, disconnect all power cords from their AC power source. For racks with a DC power distribution panel (PDP), disconnect the customer's DC power source to the PDP.

- When connecting power to the product ensure all power cables are properly connected. For racks with AC power, connect all power cords to a properly wired and grounded electrical outlet. Ensure that the outlet supplies proper voltage and phase rotation according to the system rating plate. For racks with a DC power distribution panel (PDP), connect the customer's DC power source to the PDP. Ensure that the proper polarity is used when attaching the DC power and DC power return wiring.
- Connect any equipment that will be attached to this product to properly wired outlets.
- When possible, use one hand only to connect or disconnect signal cables.
- Never turn on any equipment when there is evidence of fire, water, or structural damage.
- Do not attempt to switch on power to the machine until all possible unsafe conditions are corrected.
- When performing a machine inspection: Assume that an electrical safety hazard is present. Perform all continuity, grounding, and power checks specified during the subsystem installation procedures to ensure that the machine meets safety requirements. Do not attempt to switch power to the machine until all possible unsafe conditions are corrected. Before you open the device covers, unless instructed otherwise in the installation and configuration procedures: Disconnect the attached AC power cords, turn off the applicable circuit breakers located in the rack power distribution panel (PDP), and disconnect any telecommunications systems, networks, and modems.
- Connect and disconnect cables as described in the following procedures when installing, moving, or opening covers on this product or attached devices.
  To Disconnect:
  1. Turn off everything (unless instructed otherwise).
  2. For AC power, remove the power cords from the outlets.
  3. For racks with a DC power distribution panel (PDP), turn off the circuit breakers located in the PDP and remove the power from the Customer's DC power source.
  4. Remove the signal cables from the connectors.
  5. Remove all cables from the devices.

  To Connect:
  1. Turn off everything (unless instructed otherwise).
  2. Attach all cables to the devices.
  3. Attach the signal cables to the connectors.
  4. For AC power, attach the power cords to the outlets.
  5. For racks with a DC power distribution panel (PDP), restore the power from the Customer's DC power source and turn on the circuit breakers located in the PDP.
  6. Turn on the devices.

Sharp edges, corners and joints may be present in and around the system. Use care when handling equipment to avoid cuts, scrapes and pinching. (D005)

DANGER: Heavy equipment — personal injury or equipment damage might result if mishandled. (D006)
DANGER: Professional movers are to be used for all relocation activities. Serious injury or death may occur if systems are handled and moved incorrectly. (D008)
DANGER: Serious injury or death can occur if loaded lift tool falls over or if a heavy load falls off the lift tool. Always completely lower the lift tool load plate and properly secure the load on the lift tool before moving or using the lift tool to lift or move an object. (D010)

Caution notices for IBM FlashSystem A9000R systems

Ensure that you understand the caution notices for IBM FlashSystem A9000R systems.

Caution notices

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

CAUTION: Only trained service personnel may replace this battery. The battery contains lithium. To avoid possible explosion, do not burn or charge the battery.
Do Not:

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

Exchange only with the approved part. Recycle or discard the battery as instructed by local regulations. (C002)
CAUTION: The doors and covers to the product are to be closed at all times except for service by trained service personnel. All covers must be replaced and doors locked at the conclusion of the service operation. (C013)
CAUTION: This product is equipped with a 3-wire (two conductors and ground) power cable and plug. Use this power cable with a properly grounded electrical outlet to avoid electrical shock. (C018)

Special caution and safety notices

This information describes special safety notices that apply to FlashSystem A9000R. These notices are in addition to the standard safety notices supplied and address specific issues relevant to the equipment provided.

- **Laser safety**
  When using an NVRAM5 or NVRAM6 cluster media converter, the storage system must be installed in a restricted access location.
- **Ladder usage**
  A step or platform ladder might be necessary to service higher modules.

- **Fire suppression systems**
  A fire suppression system is the responsibility of the customer. The insurance underwriter, local fire marshal, or a local building inspector (or all three) must be consulted in selecting a fire suppression system that provides the correct level of coverage and protection.
- **Power cables**
  Use only power cables that are IBM® approved, certified, or both.
- **Sound pressure**
  Hearing protection must be worn while you service the FlashSystem A9000R system.
- **Leakage current**
  The FlashSystem A9000R system incorporates electromagnetic interference filter capacitors that are required to prevent electrical noise from penetrating the power grid. A characteristic of filter capacitors, during normal operation, is a high amount of leakage current.
- **Site preparation**
  The IBM service representative can only minimally reposition the rack at the installation site, as needed to service the FlashSystem A9000R system. The customer is responsible for using professional movers or riggers in the case of equipment relocation or disposal.

Laser safety

When using an NVRAM5 or NVRAM6 cluster media converter, the storage system must be installed in a restricted access location.

CAUTION:
This product contains a Class 1M laser. Do not view directly with optical instruments. (C028)
This equipment contains Class 1 laser products, and complies with FDA radiation regulations 21 CFR Subchapter J, international laser safety standard IEC 60825 parts -1 and -2, and relevant national standards based on these.
CAUTION:

Data processing environments can contain equipment transmitting on system links with laser modules that operate at greater than Class 1 power levels. For this reason, never look into the end of an optical fiber cable or open receptacle. (C027)
Attention: In the United States, use only SFP or GBIC optical transceivers that comply with the FDA laser registration, reporting and accessions per the Center for Devices and Radiological Health (CDRH) according to 21 CFR Subchapter J. Internationally, use only SFP or GBIC optical transceivers that comply with IEC standard 60825–1. Optical products that do not comply with these standards might produce light that is hazardous to the eyes.
Usage restrictions: The optical ports of the modules must be terminated with an optical connector or with a dust plug.

Ladder usage

A step or platform ladder might be necessary to service higher modules.

Use an OSHA/CSA approved non-conductive step or platform ladder specified for at least a 136.4 kg (300 lb.) load capacity.

Fire suppression systems

A fire suppression system is the responsibility of the customer. The insurance underwriter, local fire marshal, or a local building inspector (or all three) must be consulted in selecting a fire suppression system that provides the correct level of coverage and protection.

IBM® designs and manufactures equipment to internal and external standards that require certain environments for reliable operation. Because IBM does not test any equipment for compatibility with fire suppression systems, IBM does not make compatibility claims of any kind nor does IBM provide recommendations on fire suppression systems.

Sound pressure

Hearing protection must be worn while you service the FlashSystem A9000R system.

Attention: Depending upon local conditions, the sound pressure might exceed 85 dB(A) during service operations. When working on the FlashSystem A9000R system while either the front or rear door is in the open position, hearing protection must be worn.
CAUTION:

Depending upon local conditions, the sound pressure might exceed 85 dB(A) during service operations. Hearing protection must be worn when you are in a room that has a FlashSystem A9000R system while either the front or rear door is open or when the front and rear doors are not installed.

Leakage current

The FlashSystem A9000R system incorporates electromagnetic interference filter capacitors that are required to prevent electrical noise from penetrating the power grid. A characteristic of filter capacitors, during normal operation, is a high amount of leakage current.

Depending on the storage configuration, this leakage current can reach 100 mA.

For the most reliable operation, do not use Ground Fault Circuit Interrupter (GFCI), Earth Leakage Circuit Breaker (ELCB), and Residual Current Circuit Breaker (RCCB) type circuit breakers with a FlashSystem A9000R system. The FlashSystem A9000R system is certified for safe operation and is compliant with IEC, EN, UL, CSA 60950-1 standards. However, if leakage detection circuit breakers are required by local electrical practice, the breakers must be sized for a leakage-current rating of 300 mA or greater to reduce the risk of server outage caused by erroneous and spurious tripping.

Site preparation

The IBM® service representative can only minimally reposition the rack at the installation site, as needed to service the FlashSystem A9000R system. The customer is responsible for using professional movers or riggers in the case of equipment relocation or disposal.

Attention: Do not tilt the FlashSystem A9000R system rack more than 10 degrees, as depicted in Figure 1.

Figure 1. Rack tilt allowance

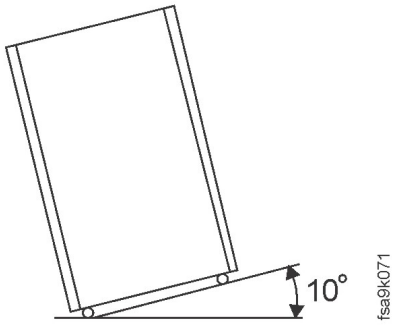

If more clearance is needed for FlashSystem A9000R or racks the Height Reduction feature code, AFR2, should be ordered to reduce the height of the rack.

When the height reduced shipping option is ordered, the doors, side panels, and rack top cover are removed before the IBM FlashSystem® A9000R is moved to its final location by professional movers. These components must then be installed during the IBM FlashSystem A9000R installation.

If tilting or rack weight reduction is needed for IBM FlashSystem A9000R shipment or delivery, the Weight Reduction feature code, AFR3, should be ordered.

When the weight reduced shipping option is ordered, grid controllers and flash enclosures are shipped separately. This allows the rack to be tilted as much as necessary in order to fit under low doorways. As a result, these components must be installed during IBM FlashSystem A9000R installation.

IBM FlashSystem A9000R arrives fully assembled with all components in place, unless the Height Reduction or Weight Reduction feature code shipping options are ordered.

DANGER

Heavy equipment - personal injury or equipment damage might result if mishandled. Use only professional movers.

Environmental notices

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

The *IBM Systems Environmental Notices* information includes statements on limitations, product information, product recycling and disposal, battery information, flat panel display, refrigeration and water-cooling systems, external power supplies, and safety data sheets.

Electromagnetic compatibility Class A notices

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

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

- **Canada Notice**
- **European Community and Morocco Notice**
- **Germany Notice**
- **Japan Electronics and Information Technology Industries Association (JEITA) Notice**
- **Japan Voluntary Control Council for Interference (VCCI) Notice**
- **Korea Notice**
- **People's Republic of China Notice**
- **Russia Notice**
- **Taiwan Notice**
- **United Kingdom Notice**
- **United States Federal Communications Commission (FCC) Notice**

Canada Notice

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

European Community and Morocco Notice

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

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

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

Germany Notice

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

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

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

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

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

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

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

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

International Business Machines Corp.
New Orchard Road

Armonk, New York 10504
Tel: 914-499-1900

Der verantwortliche Ansprechpartner des Herstellers in der EU ist:

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

Generelle Informationen:

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

Japan Electronics and Information Technology Industries Association (JEITA) Notice

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

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

高調波電流規格　JIS C 61000-3-2 適合品

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

高調波電流規格　JIS C 61000-3-2 準用品

本装置は、「高圧又は特別高圧で受電する需要家の高調波抑制対策ガイドライン」対象機器（高調波発生機器）です。
- 回路分類：6（単相、ＰＦＣ回路付）
- 換算係数：0

This statement applies to products greater than 20 A per phase, three-phase.

高調波電流規格　JIS C 61000-3-2 準用品

本装置は、「高圧又は特別高圧で受電する需要家の高調波抑制対策ガイドライン」対象機器（高調波発生機器）です。
- 回路分類：5（3相、ＰＦＣ回路付）
- 換算係数：0

Japan Voluntary Control Council for Interference (VCCI) Notice

この装置は、クラスＡ機器です。この装置を住宅環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

VCCI－A

Korea Notice

이 기기는 업무용 환경에서 사용할 목적으로 적합성평가를 받은 기기로서
가정용 환경에서 사용하는 경우 전파간섭의 우려가 있습니다.

People's Republic of China Notice

警告:在居住环境中,运行此设备可能会造成无线电干扰。

Russia Notice

ВНИМАНИЕ! Настоящее изделие относится к классу А.
В жилых помещениях оно может создавать
радиопомехи, для снижения которых необходимы
дополнительные меры

Taiwan Notice

CNS 13438

警告使用者 :
此為甲類資訊技術設備，
於居住環境中使用時，可
能會造成射頻擾動，在此
種情況下，使用者會被要
求採取某些適當的對策。

CNS 15936

警告：為避免電磁干擾，本產品不應安裝或使用於住宅環境。

IBM Taiwan Contact Information:

台灣IBM 產品服務聯絡方式：
台灣國際商業機器股份有限公司
台北市松仁路7號3樓
電話：0800-016-888

United Kingdom Notice

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

United States Federal Communications Commission (FCC) Notice

This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his own expense.

Properly shielded and grounded cables and connectors must be used in order to meet FCC emission limits. IBM® is not responsible for any radio or television interference caused by using other than recommended cables and connectors, or by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the user's authority to operate the equipment.

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions:

(1) this device might not cause harmful interference, and (2) this device must accept any interference received, including interference that might cause undesired operation.

Responsible Party:
International Business Machines Corporation
New Orchard Road
Armonk, NY 10504
Contact for FCC compliance information only: fccinfo@us.ibm.com

Electromagnetic compatibility Class B notices

The following Class B statements apply to features designated as electromagnetic compatibility (EMC) Class B in the feature installation information.

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

- **Canada Notice**
- **European Community and Morocco Notice**
- **Germany Notice**
- **Japan Electronics and Information Technology Industries Association (JEITA) Notice**
- **Japan Voluntary Control Council for Interference (VCCI) Notice**
- **Taiwan Notice**
- **United States Federal Communications Commission (FCC) Notice**

Canada Notice

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

European Community and Morocco Notice

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

Germany Notice

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

Dieses Produkt entspricht den Schutzanforderungen der EU-Richtlinie 2014/30/EU zur Angleichung der Rechtsvorschriften über die elektromagnetische Verträglichkeit in den EU-Mitgliedsstaatenund hält die Grenzwerte der EN 55022/ EN 55032 Klasse B ein.

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

Deutschland: Einhaltung des Gesetzes über die elektromagnetische Verträglichkeit von Geräten

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

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

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

Verantwortlich für die Einhaltung der EMV Vorschriften ist der Hersteller:
International Business Machines Corp.
New Orchard Road
Armonk, New York 10504
Tel: 914-499-1900

Der verantwortliche Ansprechpartner des Herstellers in der EU ist:
IBM® Deutschland GmbH
Technical Relations Europe, Abteilung M456
IBM-Allee 1, 71139 Ehningen, Germany
Tel: +49 (0) 800 225 5426
email: HalloIBM@de.ibm.com
Generelle Informationen:

Das Gerät erfüllt die Schutzanforderungen nach EN 55024 und EN 55032 Klasse B

Japan Electronics and Information Technology Industries Association (JEITA) Notice

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

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

高調波電流規格　JIS C 61000-3-2 適合品

These statements apply to products greater than 20 A, single phase.

高調波電流規格　JIS C 61000-3-2 準用品

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

- 回路分類：6（単相、ＰＦＣ回路付）
- 換算係数：0

This statement applies to products greater than 20 A per phase, three-phase.

高調波電流規格　JIS C 61000-3-2 準用品

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

- 回路分類：5（3相、ＰＦＣ回路付）
- 換算係数：0

Japan Voluntary Control Council for Interference (VCCI) Notice

この装置は、クラスB機器です。この装置は、住宅環境で使用することを目的としていますが、この装置がラジオやテレビジョン受信機に近接して使用されると、受信障害を引き起こすことがあります。

取扱説明書に従って正しい取り扱いをして下さい。　**VCCI ー B**

Taiwan Notice

IBM Taiwan Contact Information:

United States Federal Communications Commission (FCC) Notice

This equipment has been tested and found to comply with the limits for a Class B digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation. If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Connect the equipment into an outlet on a circuit different from that to which the receiver is connected.
- Consult an IBM-authorized dealer or service representative for help.

Properly shielded and grounded cables and connectors must be used in order to meet FCC emission limits. Proper cables and connectors are available from IBM-authorized dealers. IBM is not responsible for any radio or television interference caused by using other than recommended cables and connectors or by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the user's authority to operate the equipment.

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) this device might not cause harmful interference, and (2) this device must accept any interference received, including interference that might cause undesired operation.

Responsible Party:
International Business Machines Corporation
New Orchard Road
Armonk, NY 10504
Contact for FCC compliance information only: fccinfo@us.ibm.com

Release notes

Refer to the following sections for release information for the latest releases of IBM FlashSystem® A9000R, IBM® Hyper-Scale Manager, IBM XCLI Utility, and IBM Remote Support Proxy, as well as access to PDF versions of current and previous releases.

- **Release notes for IBM FlashSystem A9000R**
  The following release information is available for version 12.3.2.e of IBM FlashSystem A9000R.
- **Release notes for IBM HyperSwap Quorum Witness**
  The following release information is available for IBM Spectrum Accelerate Family HyperSwap® Quorum Witness 1.0.3.
- **Release notes for IBM Hyper-Scale Manager**
  Use IBM Hyper-Scale Manager to manage IBM Spectrum Accelerate systems.
- **Release notes for IBM XCLI Utility**
  The following release information is available for version 5.5.4.1 of IBM Extended Command-line Interface (XCLI) Utility.
- **Release notes for IBM Remote Support Proxy**
  The following release information is available for version 1.3.2.2 of IBM Remote Support Proxy.
- **Release note PDF publications**
  Use this page to access PDF versions of release notes for IBM FlashSystem A9000R, IBM Spectrum Accelerate Family HyperSwap Quorum Witness, IBM Hyper-Scale Manager, IBM Extended Command-Line Interface (XCLI) Utility, and IBM Remote Support Proxy.

Release notes for IBM FlashSystem A9000R

The following release information is available for version 12.3.2.e of IBM FlashSystem® A9000R.

What's new in 12.3.2.e

Software version 12.3.2.e is a minor software release that enhances cipher suites for SSL/TLS connections and resolves various found issues.

General availability (eGA) date: 18 January 2023

Enhanced cipher suites for SSL/TLS connections

FlashSystem A9000R 12.3.2.e updates SSL/TLS cipher suites on port 7778 (XCLI) (and on port 5989, when CIM (SMI-S server) is enabled). Elliptic Curve Diffie-Hellman Ephemeral (ECDHE) and Diffie-Hellman Ephemeral (DHE) based cipher suites in TLS/SSL connections are now supported. Additionally cipher suites using SHA1, or using anonymous DH cipher suites, or effectively using DH authentication (i.e., the certificates carry DH keys), or effectively using ECDH authentication (i.e., the certificates carry ECDH keys) are not supported by default.

The following is the full list of cipher suites that are now supported on FlashSystem A9000R storage systems on port 7778 (and on port 5989, when CIM (SMI-S server) is enabled):

- AES128-GCM-SHA256
- AES128-SHA256
- AES256-GCM-SHA384
- AES256-SHA256
- DHE-DSS-AES128-GCM-SHA256 (new)
- DHE-RSA-AES128-GCM-SHA256 (new)
- DHE-RSA-AES128-SHA256 (new)
- DHE-DSS-AES128-SHA256 (new)
- DHE-DSS-AES256-GCM-SHA384 (new)
- DHE-RSA-AES256-GCM-SHA384 (new)
- DHE-RSA-AES256-SHA256 (new)
- DHE-DSS-AES256-SHA256 (new)
- ECDHE-ECDSA-AES128-GCM-SHA256 (new)
- ECDHE-ECDSA-AES128-SHA256 (new)
- ECDHE-ECDSA-AES256-GCM-SHA384 (new)
- ECDHE-ECDSA-AES256-SHA384 (new)
- ECDHE-RSA-AES128-GCM-SHA256 (new)
- ECDHE-RSA-AES256-GCM-SHA384 (new)
- ECDHE-RSA-AES128-SHA256 (new)
- ECDHE-RSA-AES256-SHA384 (new)

Note: Enable TLS v1.2 using the **protocol_config_set** CLI command. For more information see Setting configuration parameters for a communication protocol.

Miscellaneous resolved issues

For information about the resolved issues in version 12.3.2.e, see 12.3.2.e (January 2023).

Hot upgrade from earlier versions

Software version 12.3.2.e supports direct non-disruptive upgrade (hot upgrade) of IBM FlashSystem® A9000R from the following versions.

- 12.3.2, 12.3.2.a, 12.3.2.a-1, 12.3.2.a-3, 12.3.2.b, 12.3.2.c, 12.3.2.c-1, 12.3.2.d
- 12.3.1
- 12.2.1.b-1, 12.2.1.b-2

Note: These hot upgrade paths are valid for the time of the general availability (eGA) date of version 12.3.2.e. However, hot upgrade from additional versions could be supported later on. For up-to-date information, contact IBM® Support.
The upgrade procedure must be performed by an authorized IBM service technician.

Prior to initiating software upgrade, consult with IBM Support regarding the required time allocation and expected duration of the upgrade procedure.
Note: The system e-license may automatically prompt in the management GUI upon the first post-upgrade operation, if the e-license was updated since it was last accepted. In such a case, the new e-license must be accepted before any GUI operation can be performed.

Management and solutions

FlashSystem A9000R 12.3.2.e can be managed and monitored with IBM® Hyper-Scale Manager 5.5.2 or later.

Hyper-Scale Manager 5.5.2 provides an advanced web-based graphical user interface (GUI) from which one or more IBM Spectrum Accelerate Family systems can be managed and monitored in real time from a web browser.
Note: Versions of Hyper-Scale Manager earlier than 5.5.2 are not fully supported and must be upgraded to 5.5.2 or later.
FlashSystem A9000R can also be managed from its command-line interface (CLI) or through representational state transfer (REST) application programming interfaces (APIs).

FlashSystem A9000R customers are entitled to free use of Hyper-Scale Manager, which can be obtained at any time from the IBM Fix Central website .

IBM Spectrum Accelerate Family HyperSwap Quorum Witness

The HyperSwap® functionality feature in FlashSystem A9000R 12.3.2.e is supported by IBM Spectrum Accelerate Family HyperSwap Quorum Witness 1.0.2 or later.

To significantly improve the transparent failover and facilitate the constant coordination between two IBM FlashSystem® A9000 or A9000R storage systems in a HyperSwap solution, an independent Quorum Witness software component is preferably installed at a separate site (third failure domain).
Important: Quorum Witness cannot reside on the same set of FlashSystem A9000 and/or A9000R systems of which it serves as part of a HyperSwap configuration.

Note: Newer versions of the Quorum Witness software may be released independently to better support or enhance the HyperSwap operation. For more information, see the latest Quorum Witness release notes and user guide.

IBM block storage CSI driver

FlashSystem A9000R 12.3.2.e is supported by IBM block storage CSI driver 1.0.0 and higher. The CSI (Container Storage Interface) driver is leveraged by Kubernetes persistent volumes (PVs) to dynamically provision for block storage used with stateful containers. IBM block storage CSI driver is based on an open-source IBM project (CSI driver), included as a part of IBM storage orchestration for containers.

The CSI driver can be downloaded, installed, and configured through the Red Hat® OperatorHub or from Docker Hub, through a command line terminal.

For more information, see IBM block storage CSI driver documentation.

IBM Spectrum Connect

FlashSystem A9000R 12.3.2.e is supported by IBM Spectrum® Connect 3.6.0 or later, which provides the following platform integration tools.

- IBM Storage Provider for VMware VASA
- IBM Storage Enhancements for VMware vSphere Web Client
- IBM Storage Plug-in for VMware vRealize Orchestrator
- IBM Storage Management Pack for VMware vRealize Operations Manager
- IBM Storage Automation Plug-in for Microsoft PowerShell
- IBM Storage Enabler for Containers

For more information, see IBM Spectrum Connect documentation , and the IBM Spectrum Connect marketing page .

FlashSystem A9000R customers are entitled to free use of Spectrum Connect, which can be obtained at any time from either the Hyper-Scale Manager software package or from the IBM Fix Central website .

Available cloud storage solutions

FlashSystem A9000R 12.3.2.e can be integrated with various independent software vendor (ISV) platform, application, virtualization, and cloud environments.

To facilitate this integration, IBM provides the following software solutions, also referred to as cloud storage solutions.

Cloud storage solution	Version	Compatibility note
IBM Storage Driver for OpenStack	2.3.0 or later	For use with OpenStack Cinder (OpenStack Block Storage) nodes
IBM Storage Management Pack for Microsoft System Center Operations Manager (SCOM)	2.7.0 or later	
IBM Spectrum Accelerate Family Provider for Microsoft Windows Volume Shadow Copy Service (VSS)	2.10.0 or later	
IBM Spectrum Accelerate Family Storage Replication Adapter (SRA)	3.0.1 or later	For use with VMware Site Recovery Manager 6.x or later. Requires the use of synchronous or asynchronous remote mirroring.

FlashSystem A9000R customers are entitled to free use of the IBM cloud storage solutions, which can be obtained at any time from the IBM Fix Central website.

For more information about the available cloud storage solutions for FlashSystem A9000R, see Platform and application integration..

Change log

This change log summarizes the enhancements and fixes in the different 12.3.x code level versions of IBM FlashSystem® A9000R.

Note: New functional features of these versions are only briefly summarized. For a detailed summary of the new functional features of a specific version, refer to the 'What's new' section of its release notes.

- **12.3.2.e (January 2023)**
  Software version 12.3.2.e resolves the following issues.
- **12.3.2.d (March 2022)**
  Software version 12.3.2.d further enhances internal interconnect network monitoring and resolves the following issues.
- **12.3.2.c (January 2021)**
  Software version 12.3.2.c further enhanced internal interconnect network monitoring and resolved the following issues.
- **12.3.2.b (May 2020)**
  Software version 12.3.2.b enhanced internal interconnect network monitoring and resolved the following issues.
- **12.3.2.a-3 (May 2020)**
  Software version 12.3.2.a-3 resolved the following issues.
- **12.3.2.a-2 (December 2019)**
  Software version 12.3.2.a-2 resolved the following issue.
- **12.3.2.a-1 (September 2019)**
  Software version 12.3.2.a-1 resolved the following issue.
- **12.3.2.a (August 2019)**
  Software version 12.3.2.a resolved the following issues.
- **12.3.2 (May 2019)**
  Software version 12.3.2 added support for virtual LAN (VLAN) capabilities, and resolved the following issues.
- **12.3.1 (January 2019)**
  Software version 12.3.1 of IBM FlashSystem A9000R introduced deduplication-aware capacity management (available via IBM® Hyper-Scale Manager 5.5.1

or later) and enabled the expansion of the new minimal rack configuration that was introduced with version 12.3.0..

- **12.3.0.a (November 2018)**
  Software version 12.3.0.a of IBM FlashSystem A9000R resolved the following issues.
- **12.3.0 (November 2018)**
  Software version 12.3.0 of IBM FlashSystem A9000R added the following functional features.

12.3.2.e (January 2023)

Software version 12.3.2.e resolves the following issues.

Ticket ID	Description
SYS-317567	**Fixed:** In a HyperSwap relation, where system A has a Primary role, and system B has a Secondary role, in rare cases, due to a microcode defect in the health monitoring mechanism on system A, the system might incorrectly set one of the internal parameters. If system B later undergoes a HyperSwap failover, the HyperSwap volumes on system A might not be properly 'blocked' for I/O and not made Secondary. This can result in a Primary-Primary condition that can cause an impact to the data. This data impact can range from no impact to, in very rare cases, undetected data corruption. Severity: HIPER Affected versions: 12.1.0 or later
SYS-316999	**Fixed:** Security scans may report SSH server (host) RSA public keys for ports 22 and 2222 are smaller than recommended. Severity: High Impact Affected versions: 12.0.0 or later
SYS-317109	**Fixed:** When using external plugins or applications that depend on the output of the **mirror_list** command to identify the current volume or consistency group name, renaming a mirrored volume or consistency group (using the **vol_rename** / **cg_rename** command), can temporarily cause inaccessibility by the given plugin application. During this time, the **mirror_list** output does not reflect the new volume/consistency group name, while the **vol_list** / **cg_list** output reflects the name correctly. In certain rare cases, this temporary discrepancy can lead to discrepancies during failover operations for the mirrored volume by the external tools, affecting their functionality. Severity: High Impact Affected versions: 12.0.0 or later
SYS-317173	**Fixed:** Common security issues. For more information about these security issues, see the following on the Common Vulnerabilities and Exposures (CVE) information website (cve.org): CVE-2017-17807, CVE-2017-18595, CVE-2018-7191, CVE-2018-19985, CVE-2018-20169, CVE-2018-25032, CVE-2018-20856, CVE-2018-25032, CVE-2019-3901, CVE-2019-11190, CVE-2019-12382, CVE-2019-13648, CVE-2019-14821, CVE-2019-15916, CVE-2020-10711, CVE-2020-25709, CVE-2020-25710, CVE-2020-36322, CVE-2021-3733 (Python 2), CVE-2021-3737 (Python 2), CVE-2021-4189, CVE-2021-20271, CVE-2021-27219, CVE-2021-28831, CVE-2021-28971, CVE-2021-37750, CVE-2021-41617, CVE-2021-43527, CVE-2021-45485, CVE-2021-45960, CVE-2021-46143, CVE-2022-0391, CVE-2022-0778, CVE-2022-22822, CVE-2022-22823, CVE-2022-22824, CVE-2022-22825, CVE-2022-22826, CVE-2022-22827, CVE-2022-23852, CVE-2022-24407, CVE-2022-25235, CVE-2022-25236, CVE-2022-25315 Severity: High Impact Affected versions: 12.0.0 or later
SYS-317442	**Fixed:** Common security issues. For more information about these security issues, see the following on the Common Vulnerabilities and Exposures (CVE) information website (cve.org): CVE-2017-6001, CVE-2019-20811, CVE-2020-28097, CVE-2020-36557, CVE-2020-36558, CVE-2021-3347, CVE-2021-20322, CVE-2022-0492, CVE-2022-0850, CVE-2022-1012, CVE-2022-1271, CVE-2022-1292, CVE-2022-1353, CVE-2022-1729, CVE-2022-2068, CVE-2022-22942, CVE-2022-36946 Severity: High Impact Affected versions: 12.0.0 or later
SYS-317566	**Fixed:** In rare cases, during a HyperSwap failover, when a volume is promoted from Secondary to Primary role and the relations pass to the correct system, the system might incorrectly fail software component(s). This can occur where a system is in a HyperSwap relation and the system has HyperSwap relations that are from both Primary and Secondary roles. In extremely rare cases, this can lead to loss of access (either partial or full) to the system. This issue also affects systems using Hyper-Scale Mobility during delete operations with both Hyper-Scale Mobility proxy and owner volumes present on the same system. Severity: High Impact Affected versions: 12.1.0 or later
SYS-317148	**Fixed:** In rare cases, certain indeterminate hardware issues on a grid controller might cause issues on the internal InfiniBand network, resulting in a failure to restart internal network interface(s). This could affect system internal communication and lead to loss of access. Severity: Moderate Affected versions: 12.3.2.c or later

12.3.2.d (March 2022)

Software version 12.3.2.d further enhances internal interconnect network monitoring and resolves the following issues.

Ticket ID	Description
SYS-316394	**Fixed:** In some cases, if the host internally uses the infrequently used **request_sense** SCSI command, this may conflict with the Hyper-Scale mobility at the proxy phase, potentially causing interface node failures. In extremely rare cases, this might also occur within HyperSwap® relations, also potentially resulting in interface node failures. Severity: HIPER Affected versions: 12.0.0 or later
SYS-316571	**Fixed:** When mirror offline initialization is used and the volume mirror is put into a consistency group before the initialization of that volume mirror has completed, in some rare cases the secondary volume on the remote system is not properly synchronized. In these instances, one of the following may occur: • In cases where a host erroneously reads from an unallocated area of the disk (on a remote copy of the volume) and the host reads a partition that is allocated only on the secondary (and not on the primary), the host might read stale data, resulting in undetected data corruption. • If a host writes to a part of a partition on a primary volume copy, and the partition at hand was allocated only on the secondary volume copy, the system will detect the data discrepancy and will emit the `SCRUBBING_REMOTE_DIGEST_DIFF` event. If this occurs the automatic scrubbing cycle that periodically runs on the system identifies and fixes this inconsistency. Severity: HIPER Affected versions: 12.3.0 or later
SYS-316084	**Fixed:** In rare cases, internal disconnections within the flash enclosure may lead to loss of access. Severity: High Impact Affected versions: 12.0.0 or later
SYS-316240	**Fixed:** In rare cases, an automatic error handling correction in the flash enclosure may not fully clear the error state. If an additional error occurs, it may not be properly handled. As a result, the flash enclosure may become unreachable, causing an emergency shutdown. Severity: High Impact Affected versions: 12.1.0 or later
SYS-316472	**Fixed:** In extremely rare cases, local metadata describing zeros might become corrupted, leading to erroneous data being written or returned to the host instead of the zeros, potentially leading to loss of access. Severity: High Impact Affected versions: 12.0.0 or later
SYS-316635	**Fixed:** In some cases, when a MicroLatency module experiences excessive response times it may not be marked as failed. This leads to timeouts and performance degradation that may lead to loss of access. Severity: High Impact Affected versions: 12.0.0 or later
SYS-316693	**Fixed:** Common security issues. For more information about these security issues, see the following on the Common Vulnerabilities and Exposures (CVE) information website (cve.org): CVE-2014-3591, CVE-2014-4617, CVE-2016-4658, CVE-2018-7755, CVE-2018-9516, CVE-2018-9517, CVE-2018-11574, CVE-2018-13093, CVE-2018-13094, CVE-2018-13095, CVE-2018-14734, CVE-2018-16658, CVE-2018-16884, CVE-2018-18281, CVE-2018-20852, CVE-2019-2974, CVE-2019-3819, CVE-2019-3882, CVE-2019-3892, CVE-2019-5489, CVE-2019-7222, CVE-2019-11599, CVE-2019-11810, CVE-2019-11833, CVE-2019-12450, CVE-2019-14835, CVE-2019-18348, CVE-2019-19532, CVE-2019-19768, CVE-2019-20386, CVE-2020-1971, CVE-2020-2574, CVE-2020-2752, CVE-2020-2780, CVE-2020-2812, CVE-2020-8492, CVE-2020-8597, CVE-2020-8625, CVE-2020-10543, CVE-2020-10769, CVE-2020-10878, CVE-2020-12243, CVE-2020-12723, CVE-2020-14314, CVE-2020-14331, CVE-2020-14422, CVE-2020-15436, CVE-2020-15862, CVE-2020-25211, CVE-2020-25212, CVE-2020-25648, CVE-2020-25656, CVE-2020-25668, CVE-2020-25684, CVE-2020-25685, CVE-2020-25686, CVE-2020-25692, CVE-2020-26116, CVE-2020-27619, CVE-2020-29370, CVE-2020-29660, CVE-2021-3177, CVE-2021-3426, CVE-2021-3712, CVE-2021-3733, CVE-2021-3737, CVE-2021-20261 , CVE-2021-20265, CVE-2021-20305, CVE-2021-23336, CVE-2021-23840, CVE-2021-23841, CVE-2021-25214, CVE-2021-25217, CVE-2021-26937, CVE-2021-27363, CVE-2021-27364, CVE-2021-27365, CVE-2021-29154, CVE-2021-29650, CVE-2021-33909 Severity: High Impact Affected versions: 12.0.0 or later
SYS-316695	**Fixed:** In rare cases where a flash enclosure is in a degraded state following two MicroLatency module failures, one of which was the spare MicroLatency module, and the other MicroLatency module has been revived and the flash enclosure rebuild completed, potential loss of access may occur. Severity: High Impact Affected versions: 12.3.2.c or later
SYS-316803	**Fixed:** In rare cases, during a failed grid controller automatic isolation following its failure might cause an incomplete HyperSwap failover, causing an active-active pairing, potentially leading to data loss. Severity: High Impact Affected versions: 12.1.0 or later

Ticket ID	Description
SYS-316905	**Fixed:** In rare cases, a system that is in HyperSwap relation may experience performance degradation and potential loss of access due to an internal network congestion that exists on a high-availability (HA) peer system. This could occur due to a HyperSwap error in detecting the condition of the HA peer system. Severity: High Impact Affected versions: 12.1.0 or later
SYS-316292	**Fixed:** In rare cases, when ACA mode is used on the host side as part of the SCSI protocol (most commonly used by AIX® hosts), the ACA condition causes the interface node to incorrectly stop all external I/O operations. This may potentially lead to, among other scenarios, one or more interface node failures that can lead to loss of access or loss of synchronization in synchronous mirror relationships. Severity: Moderate Affected versions: 12.0.0 or later
SYS-316335	**Fixed:** When a user's LDAP group membership role is located beyond the 100th location in the membership list, authentication to that role is failed. In these cases, login and tasks are blocked and the `LOGIN_FAILURE_USER_HAS_NO_RECOGNIZED_ROLE` error code is returned. Note: The tasks that are blocked are user role dependent. Severity: Moderate Affected versions: 12.0.0 or later
SYS-316370	**Fixed:** In an extremely rare combination of circumstances, a configuration update across all nodes that coincides with a very specific SCSI I/O operation, causes an interface node failure. Severity: Moderate Affected versions: 12.2.1 or later
SYS-316773	**Fixed:** On rare occasions, due to repeated failures of forked processes, invoked by internal communication tasks, the processes may retain freed resources leading to an out of memory condition on the grid controller. In extremely rare cases, this may lead to loss of access. Severity: Moderate Affected versions: 12.0.2 or later
SYS-316587	**Fixed:** During a flash canister replacement the new flash canister does not properly clear the old registers value for the power manager object management. As a result, the **system_average_power_prepare** and **system_average_power_consumption** commands do not return any values, necessary for tracking and planning system power consumption. Severity: Service Affected versions: 12.0.0 or later
SYS-316770	**Fixed:** InfiniBand Switch `IB_SWITCH_MISSING` event is not emitting sufficient alerts. Severity: Low Affected versions: 12.0.0 or later

12.3.2.c (January 2021)

Software version 12.3.2.c further enhanced internal interconnect network monitoring and resolved the following issues.

Ticket ID	Description
SYS-316012	**Fixed:** In extremely rare cases, multiple cache failures may occur due to internal partition table handling, resulting in loss of access. Severity: HIPER Affected versions: 12.0.0 or later
SYS-315916	**Fixed:** In rare cases, some internal asynchronous mirroring-related operations might lead to multiple cache failures, causing an emergency shutdown, leading to loss of access. Severity: High Impact Affected versions: 12.0.2 or later
SYS-315956	**Fixed:** In extremely rare cases, the system may become locked for read-only operations, due to an incorrect `SYSTEM_OUT_OF_PHYSICAL_SPACE` reading, even though capacity list shows available capacity and without previous warnings. Severity: High Impact Affected versions: 12.2.1 or later
SYS-315978	**Fixed:** (Model 415 only) In rare cases, a unique sequence of events can trigger a MicroLatency module failure. When coupled with an additional MicroLatency module failure, this can lead to the flash enclosure to go offline and loss of access. Severity: High Impact Affected versions: 12.0.0 or later
SYS-316006	**Fixed:** In very rare situations, when using asynchronous mirroring, a metadata inspection for a deleted volume may lead to loss of access. Severity: High Impact Affected versions: 12.2.0 or later
SYS-316009	**Fixed:** In very rare cases, Multi-site relation deletions (**multisite_delete**) may result in multiple cache failures, causing loss of access. Severity: High Impact Affected versions: 12.3.0 or later

Ticket ID	Description
SYS-316039	**Fixed**: In rare cases, Fibre Channel port firmware under heavy load might hang when reaching its internal queue limit. As a result, some hosts may lose access. Severity: High Impact Affected versions: 12.0.3.b or later
SYS-316045	**Fixed:** In rare cases, a transient error in the flash enclosure may result in inconsistent data to be sent to the FlashSystem A9000R system, resulting in loss of access. Severity: High Impact Affected versions: 12.0.0 or later
SYS-316069	**Fixed:** Common security issue. For more information about these security issues, see the following on the Common Vulnerabilities and Exposures (CVE) information website (cve.mitre.org): CVE-2015-2716, CVE-2015-8035, CVE-2016-5131, CVE-2017-15412, CVE-2017-18258, CVE-2018-10360, CVE-2018-14404, CVE-2018-14567, CVE-2018-18066, CVE-2018-19519, CVE-2018-20843, CVE-2018-20852, CVE-2018-5745, CVE-2019-11068, CVE-2019-11719, CVE-2019-11727, CVE-2019-11756, CVE-2019-12749, CVE-2019-14834, CVE-2019-14866, CVE-2019-14898, CVE-2019-15903, CVE-2019-16056, CVE-2019-16935, CVE-2019-17006, CVE-2019-17023, CVE-2019-17498, CVE-2019-18197, CVE-2019-19062, CVE-2019-19126, CVE-2019-19447, CVE-2019-19527, CVE-2019-19530, CVE-2019-19537, CVE-2019-19767, CVE-2019-19956, CVE-2019-20388, CVE-2019-20636, CVE-2019-20907, CVE-2019-2737, CVE-2019-2739, CVE-2019-2740, CVE-2019-2805, CVE-2019-5094, CVE-2019-5188, CVE-2019-5436, CVE-2019-5482, CVE-2019-5544, CVE-2019-6465, CVE-2019-6477, CVE-2019-9924, CVE-2020-10531, CVE-2020-10720, CVE-2020-10732, CVE-2020-10741, CVE-2020-11565, CVE-2020-11884, CVE-2020-12400, CVE-2020-12401, CVE-2020-12402, CVE-2020-12403, CVE-2020-12464, CVE-2020-12465, CVE-2020-12770, CVE-2020-12826, CVE-2020-2732, CVE-2020-6829, CVE-2020-7595, CVE-2020-8177, CVE-2020-8622, CVE-2020-8623, CVE-2020-8624, CVE-2020-8647, CVE-2020-8649, CVE-2020-8834, CVE-2020-9383, CVE-2020-9391 Severity: High Impact Affected versions: 12.0.0 or later
SYS-316182	**Fixed:** In extremely rare cases, in a system with any type of mirroring relation, a certain sequence of events, combined with a specific write pattern of "zeros" to a deactivated mirror relation, followed by a role change and specific time frame of the mirror re-activation, an internal race condition might lead to undetected data corruption on the primary system. This may cause invalid data to be returned to the host as undetected data corruption. Severity: High Impact Affected versions: 12.3.2 or later
SYS-316205	**Fixed:** In extremely rare cases, when the write_same of "zeros" is used in a combination with the **vol_copy** command, an incorrect handling of the written zeros on the backup cache may occur. If the system later performs an automatic data redistribution, this may result in one of the following: • Undetected data corruption when reading the data. • Multiple cache failures leading to loss of access during write attempts to the data. Severity: High Impact Affected versions: 12.3.0 or later
SYS-315974	**Fixed:** In rare cases, a Fibre Channel SCSI initiator may send invalid I/O packets to the Fibre Channel SCSI target port. Due to a secondary issue, this can result in SCSI port disconnects on the system. Severity: Moderate Affected versions: 12.0.0 or later
SYS-316043	**Fixed**: In extremely rare cases a request to duplicate a snapshot group may result in an internal system software component restart. This causes no impact to the system or connected hosts. Severity: Low Affected versions: 12.0.0 or later

12.3.2.b (May 2020)

Software version 12.3.2.b enhanced internal interconnect network monitoring and resolved the following issues.

Ticket ID	Description
SYS-307853	**Fixed**: In extremely rare cases, when utilizing large amounts of mirroring pairs,snapshots, random write patterns, and heavy loads, internal resource depletion may cause multiple system component failures which could potentially lead to loss of access. Severity: High Impact Affected versions: 12.0.0 or later
SYS-308559	**Fixed**: Common security issue. For more information about these security issues, see CVE-2017-15906 on the Common Vulnerabilities and Exposures (CVE) information website (cve.mitre.org). Severity: High Impact Affected versions: 12.2.1 or later
SYS-310434	**Fixed**: In rare cases, an automatic recovery of the internal interconnect network will fail, causing possible loss of access. Severity: High Impact Affected versions: 12.0.0 or later

Ticket ID	Description
SYS-315321	**Fixed**: On extremely rare occasions, the internal background process in charge of the data consistency verification might not complete it's scrubbing cycle. Severity: High Impact Affected versions: 12.1.0 or later
SYS-315429	**Fixed**: In very rare cases, a medium error from the flash enclosure may not be properly handled by the FlashSystem A9000R code, resulting in loss of access. Severity: High Impact Affected versions: 12.0.0 or later
SYS-315539	**Fixed**: Using mirrored snapshot overwrite may cause internal system software component failure, causing loss of access and, in rare cases, may also cause data loss. Severity: High Impact Affected versions: 12.0.0 or later
SYS-315555	**Fixed**: Invalid health readings may lead to MicroLatency module failures. Severity: High Impact Affected versions: 12.2.0 or later
SYS-315556	**Fixed**: In rare cases, error recovery of a flash enclosure during a loss of access event may result in minor data loss. Severity: High Impact Affected versions: 12.0.0 or later
SYS-315671	**Fixed**: Common security issue. For more information about these security issues, see the following on the Common Vulnerabilities and Exposures (CVE) information website (cve.mitre.org): CVE-2018-0495, CVE-2018-0734, CVE-2018-1122, CVE-2018-12404, CVE-2018-14647, CVE-2018-1547, CVE-2018-15686, CVE-2018-16062, CVE-2018-16402, CVE-2018-16403, CVE-2018-16842, CVE-2018-16866, CVE-2018-16888, CVE-2018-18310, CVE-2018-18520, CVE-2018-18521, CVE-2018-3058, CVE-2018-3063, CVE-2018-3066, CVE-2018-3081, CVE-2018-3282, CVE-2018-5407, CVE-2019-11729, CVE-2019-11745, CVE-2019-1559, CVE-2019-15807, CVE-2019-15807, CVE-2019-16089, CVE-2019-2503, CVE-2019-2529, CVE-2019-2614, CVE-2019-2627, CVE-2019-3855, CVE-2019-3856, CVE-2019-3857, CVE-2019-3858, CVE-2019-3861, CVE-2019-3862, CVE-2019-3863, CVE-2019-5010, CVE-2019-6454, CVE-2019-7149, CVE-2019-7150, CVE-2019-7664, CVE-2019-7665, CVE-2019-9740, CVE-2019-9947, CVE-2019-9948 Severity: High Impact Affected versions: 12.0.0 or later
SYS-315365	**Fixed**: In some cases, live vaulting, the automatic process of cache and metadata saving to the vault devices, might not succeed during heavy system loads. When this occurs, the `LIVE_VAULTING_PROCESS_CANCELLED_TOO_MANY_TIMES` event is emitted. Severity: Moderate Affected versions: 12.1.0.e or later
SYS-315411	**Fixed**: In cases where the system fails to complete the live vaulting process, an event is not emitted. Severity: Moderate Affected versions: 12.0.0 or later
SYS-315425	**Fixed**: Full Multi-site relation allowances are not able to be utilized. Depending on the scenario, one of the following events may be emitted: `MAX_SYNC_MIRRORS_REACHED`, `MAX_ASYNC_MIRRORS_REACHED`, or `MAX_MULTISITES_REACHED`. Severity: Moderate Affected versions: 12.3.0 or later
SYS-315450	**Fixed**: When an LDAP authenticated user tries listing domain-related storage configuration objects (**vol_list**, **pool_list**, etc.) using XCLI or HSM, reply will be empty. Rerunning the command (or running any other command) within 20 minutes lists the correct configuration. Severity: Service Affected versions: 12.0.0 or later
SYS-314323	**Fixed**: When running the **ldap_config_set** command in order to set the Storage Integration Administrator (SIA) role, the new SIA role does not take effect, even though the command shows that it completed successfully. Severity: Low Affected versions: 12.3.2.a or later
SYS-315632	**Fixed**: In some cases, an unnecessary `TARGET_CONNECTION_HA_SUFFICIENT` and/or `TARGET_CONNECTION_HA_INSUFFICIENT` event may be emitted when no HyperSwap® is configured. Severity: Low Affected versions: 12.1.0 or later

12.3.2.a-3 (May 2020)

Software version 12.3.2.a-3 resolved the following issues.

Ticket ID	Description

Ticket ID	Description
SYS-315562	**Fixed**: (Models 425 or U25 only) A MicroLatency module failure may result in I/O latency between the flash enclosure and the rest of the system. As a result, the host I/O cannot serve for a period of 3 - 4 minutes, with I/Os resuming automatically after MicroLatency module failure is properly identified. Severity: HIPER Affected versions: 12.3.2.a or later

12.3.2.a-2 (December 2019)

Software version 12.3.2.a-2 resolved the following issue.

Ticket ID	Description
SYS-315392	**Fixed**: The flash enclosure interface controllers may cause system outages due to internal hardware failures. Severity: High Impact Affected versions: 12.0.0 or later

12.3.2.a-1 (September 2019)

Software version 12.3.2.a-1 resolved the following issue.

Ticket ID	Description
SYS-315292	**Fixed**: System models 415 that were encrypted when on system version prior to 12.2.0 and were shutdown due to a thermal issue or manual procedure might not be able to be moved into a fully operational state without IBM® support involvement. Severity: High Impact Affected versions: 12.2.0 or later

12.3.2.a (August 2019)

Software version 12.3.2.a resolved the following issues.

Ticket ID	Description
SYS-309772	**Fixed**: Common security issue. For more information about this security issue, see CVE-2017-17833 on the Common Vulnerabilities and Exposures (CVE) information website (cve.mitre.org). Severity: High Impact Affected versions: 12.0.0 or later
SYS-313717	**Fixed**: In rare cases, a MicroLatency module communication error could sometimes cause loss of access due to a rare timing condition, resulting in a graceful emergency shutdown of the system. Severity: High Impact Affected versions: 12.0.0 or later
SYS-315184	**Fixed:** In rare cases, in a Multi-site HA DR relation, a required failover will not occur due to a problem with the primary system, following a disconnected service reporting an incorrect peer health status. Severity: High Impact Affected versions: 12.1.0 or later
SYS-314004	**Fixed**: In extremely rare cases, when the system is undergoing heavy SCSI unmap loads combined with excessive write operations, write performance will be affected. The system resumes normal performance once the heavy unmap operations complete. Severity: High Impact Affected versions: 12.0.0 or later
SYS-315076	**Fixed**: In extremely rare cases, a switch role during a temporary disconnect between the primary system and Quorum Witness, in addition to a unique sequence of events, a failover, if required, may not be able to be performed. An example of the type of sequence of events could be a connectivity issue between the primary and secondary volumes in a Multi-Site HA DR relation occurring prior to the above temporary disconnection from the Quorum Witness and the unique sequence of events. Severity: High Impact Affected versions: 12.1.0 or later
SYS-315195	**Fixed:** Common security issue. For more information about these security issues, see CVE-2019-11477 , CVE-2019-11478, and CVE-2019-11479 on the Common Vulnerabilities and Exposures (CVE) information website (cve.mitre.org). Severity: High Impact Affected versions: 12.0.0 or later

Ticket ID	Description
SYS-310286	**Fixed:** If the storage system has reached its maximum storage capacity, performance for user reads and remote mirroring might be degraded. Severity: Moderate Affected versions: 12.0.0 or later
SYS-310912	**Fixed:** When using the **host_remove_port** CLI command on a port designated for a HyperSwap® volume, the command should remove the port on both the local and the remote systems. Instead, only the local system port is removed and the virtual port on the remote system remains configured for the HyperSwap volume. This can also expose the volume to SCSI reservation conflict until the virtual port is manually removed. Severity: Moderate Affected versions: 12.1.0 or later
SYS-310952	**Fixed:** In rare cases, deleting an asynchronous mirror relation that is part of a sync job batch immediately after start, causes other sync jobs to the same remote target to get stuck. These sync job prolonged delays can be seen by using the **sync_job_list** CLI command. Severity: Moderate Affected versions: 12.0.0 or later
SYS-311218	**Fixed:** When a volume has more than 10 user metadata items, the initial sync occurring after the creation of an IBM® Hyper-Scale Mobility or HyperSwap relation may fail and management operations may not run on the system. This can be seen when performing various CLI commands and receiving the `MCL_NETWORK_ERROR` return code. This does not impact any I/O operations. Note: Note: In order to check how many metadata items are listed on a specific volume, use the **metadata_list** CLI command. Severity: Moderate Affected versions: 12.0.0 or later
SYS-314736	**Fixed:** In rare cases, in a Multi-site relation environment which includes many mirroring relations, a target reconnection may cause some I/O delays. Severity: Moderate Affected versions: 12.1.0 or later
SYS-314991	**Fixed:** In rare cases following specific software failures on a secondary system, when using synchronous mirroring and the volume is part of a consistency group, management operations on the primary system will not run. Severity: Moderate Affected versions: 12.2.1 or later
SYS-315015	**Fixed:** In some cases, due to an issue in the internal communication, a faulty `FLASH_CANISTER_ETH_LINK_MISWIRE` event is emitted. After approximately one hour, this event is followed by the correcting `FLASH_CANISTER_ETH_WIRING_OK` event. Severity: Moderate Affected versions: 12.0.0 or later
SYS-315029	**Fixed:** In a Multi-site HA/DR configuration, a user will not be able to add a Multi-site volume into a consistency group, if the following order of events have occurred: 1. A standby replication relation was activated for a specific consistency group. 2. A role change occurred. 3. The new relation is not yet been activated for this consistency group. Severity: Moderate Affected versions: 12.3.0 or later
SYS-315146	**Fixed**: Flash enclosure daily snap collections may result in momentary latency increase. Severity: Moderate Affected versions: 12.3.2.a
SYS-310343	**Fixed:** The status of the backup battery unit (BBU) in the flash enclosure might be shown as failed while it is still operating normally. If this occurs, contact IBM Support. Severity: Service Affected versions: 12.3.0 or later
SYS-314453	**Fixed:** In rare cases, latency spikes in system statistics might unnecessarily be emitted. This is a false identification and can be ignore. This behavior has no actual impact on system performance. Severity: Service Affected versions: 12.0.0 or later

12.3.2 (May 2019)

Software version 12.3.2 added support for virtual LAN (VLAN) capabilities, and resolved the following issues.

Ticket ID	Description
SYS-313334	**Fixed**: In rare conditions, a bug in the microcode might lead to undetected data corruption (a range of 16 MB returning old data instead of zeros), on the Secondary volume of any type of mirroring relationship: synchronous mirroring, asynchronous mirroring, HyperSwap®, or Multi-site HA DR. Severity: HIPER Affected versions: 12.0.0 or later

Ticket ID	Description
SYS-313522	**Fixed**: In rare cases, a CPU error might lead to a `SCSI MEDIUM ERROR` condition. Severity: High Impact Affected versions: 12.2.1 or later
SYS-308548	**Fixed**: In a HyperSwap relation, when a remote system target is deleted and then redefined, creating new HyperSwap relations on this target system is not possible. Severity: Moderate Affected versions: 12.1.0 or later
SYS-310148	**Fixed**: In rare cases, due to a unique chain of events resulting in an internal resource depletion, disconnected Fibre Channel hosts may not be able to automatically reconnect. Severity: Moderate Affected versions: 12.0.0 or later
SYS-310990	**Fixed**: In rare cases where an internal process fails, the reported connectivity for FC or iSCSI hosts might be wrong when running the **host_connectivity_list** and **fc_connectivity_list** CLI commands. Severity: Service Affected versions: 12.1.0 or later
SYS-311467	**Fixed**: Statistics information is not kept for target ports in mirror relations. Severity: Service Affected versions: 12.0.0 or later

12.3.1 (January 2019)

Software version 12.3.1 of IBM FlashSystem A9000R introduced deduplication-aware capacity management (available via IBM® Hyper-Scale Manager 5.5.1 or later) and enabled the expansion of the new minimal rack configuration that was introduced with version 12.3.0..

Version 12.3.1 also resolved the following issues.

Ticket ID	Description
SYS-310338	**Fixed**: In rare cases, for a system with encryption enabled: 1. During a replacement procedure of a malfunctioning MicroLatency module, the procedure might fail and the new MicroLatency module might stay offline, thus compromising the system redundancy. 2. During firmware upgrade of the flash enclosure, the procedure might halt and the flash enclosure upgrade can not complete. In both of the cases described above, a maintenance window might have to be scheduled for a complete restart of the system. Severity: HIPER Affected versions: 12.2.0 or later
SYS-309744	**Fixed**: In rare cases, if the InfiniBand interconnect fabric of an A9000R system becomes overloaded, an InfiniBand switch port might go offline for a short period and then get back online. If at the same time, another issue such as a port failure causes the grid controller to be connected to the InfiniBand fabric via a single InfiniBand link, the grid controller might fail. If at the same time additional grid controllers fail, hosts could lose access to the storage system. Severity: High Impact Affected versions: 12.0.0 or later
SYS-310091	**Fixed**: In rare cases, a flash RAID controller malfunction and a particular write error might cause MicroLatency modules on the flash enclosure to fail, leading to loss of access to the storage system. Severity: High Impact Affected versions: 12.2.0 or later
SYS-308309	**Fixed**: In rare cases, a replacement of a flash canister in the storage system might cause performance degradation. Severity: Moderate Affected versions: 12.0.0 or later
SYS-309381	**Fixed**: In rare cases, an InfiniBand interconnect issue might lead to performance degradation. Severity: Moderate Affected versions: 12.0.0 or later
SYS-311260	**Fixed**: If performance classes are used, the I/O bandwidth allocation that is associated with some performance classes might be incorrectly calculated. This could lead to incorrect performance limits for hosts or volumes. Severity: Moderate Affected versions: 12.0.0 or later
SYS-311325	**Fixed**: If the size of a written volume and its snapshots is larger than the total volume size, the **vol_list** CLI command shows only the volume total size, without the size of the volume snapshots. Severity: Service Affected versions: 12.0.0 or later

12.3.0.a (November 2018)

Software version 12.3.0.a of IBM FlashSystem A9000R resolved the following issues.

Ticket ID	Description
SYS-309775	**Fixed**: In rare cases, after software upgrade on a system of model 425, the system might experience growing latencies in I/O operations, until all hosts lose access to the storage system. Severity: HIPER Affected versions: 12.3.0
SYS-310311	**Fixed**: In rare cases, systems that are using snapshots or asynchronous remote mirroring might disconnect when a volume deletion operation is performed. In such a case, the `DATA_REDUCTION_TIER_BECOMING_OFFLINE` event is logged. Severity: High Impact Affected versions: 12.2.0 or later

12.3.0 (November 2018)

Software version 12.3.0 of IBM FlashSystem A9000R added the following functional features.

- Multi-site high availability (HA) and disaster recovery (DR).
- Removal of storage pool size limitation (was 1 PB).
- Monitor internal/external latency at the system level.
- Pending deletion indication.
- Option for using a longer user password of up to 64 characters.
- Option for acquiring a new entry-level configuration of one flash enclosure and three grid controllers, reducing the initial minimum capacity by nearly 50%. The expansion option for the new minimal rack configuration (adding grid controllers and flash enclosures) is to be enabled by a future software version.

Version 12.3.0 also resolved the following issues.

Ticket ID	Description
SYS-307098	**Fixed**: In a HyperSwap® setup with an active SCSI reservation, if only a subset of the host ports that are defined on the System holding the Primary volume, is defined on the System which holds the Secondary volume, the latter system might not conform to the SCSI reservation. As a result, the Secondary volume might be read by another host; and after a HyperSwap failover, the new Primary volume (former Secondary volume) might be written to by another host. This violation of SCSI reservation will be undetected. Severity: HIPER Affected versions: 12.1.0 or later
SYS-305091	**Fixed**: In rare cases, using the **io_pause** CLI command might cause loss of access to the storage system. Severity: High Impact Affected versions: 12.0.0 or later
SYS-305121	**Fixed**: When using IBM® Hyper-Scale Mobility and the volume migration is in Proxy state, a broken link to the destination system might cause hosts to lose access to one or more of the Proxy-state interface modules. Severity: High Impact Affected versions: 12.1.0 or later
SYS-306086	**Fixed**: When most of the physical memory of the storage system is occupied by different operation processes, deleting a significant amount of obsolete data and then restarting I/O operations before the data garbage collection process has completed - might result in loss of access of hosts to the storage system. Severity: High Impact Affected versions: 12.0.0 or later
SYS-306139	**Fixed**: Before a hot upgrade, all primary HyperSwap volumes must be deactivated in order to make the secondary volumes unavailable. Otherwise local I/O operations during the hot upgrade will be blocked. Severity: High Impact Affected versions: 12.1.0 or later
SYS-306457	**Fixed**: In rare cases, the data redistribution after the addition of a grid element might fail due to a data transfer error. If this occurs, hot upgrade of the storage system software will not be possible. Severity: High Impact Affected versions: 12.0.3 or later
SYS-306775	**Fixed**: The iSCSI port driver might crash when invalid data length is specified in a read or write I/O command, causing loss of access of all iSCSI hosts to the storage system. Severity: High Impact Affected versions: 12.0.0 or later

Ticket ID	Description
SYS-307087	**Fixed**: Common security issues. For more information about these security issues, see CVE-2016-3672, CVE-2017-9725, CVE-2017-12190, CVE-2017-14140, CVE-2017-15126, CVE-2017-17449, CVE-2017-17558, CVE-2017-18017, and CVE-2018-6927 on the Common Vulnerabilities and Exposures (CVE) information website (cve.mitre.org). Severity: High Impact Affected versions: 12.0.0 or later
SYS-307118	**Fixed**: In rare cases, a HyperSwap relation might get disconnected when trying to synchronize SCSI registrations between the Primary and the Secondary systems. Severity: High Impact Affected versions: 12.1.0 or later
SYS-307626	**Fixed**: In rare cases, a storage system might fail to isolate a failing component, but wrongly detect that the component is successfully isolated. This might result in hosts losing access to the storage system, or in data unavailability. Severity: High Impact Affected versions: 12.0.0.a or later
SYS-307692	**Fixed**: In extremely rare cases, a grid controller might not be able to exchange data with one of the flash enclosures, leading to a severe impact on performance, and possibly also to loss of access to the storage system. Severity: High Impact Affected versions: 12.0.0.a or later
SYS-307426	**Fixed**: In a HyperSwap relation, a host issuing a LUN reset task might cause failure of all subsequent reservation operations, causing the HyperSwap relation to break. Severity: High Impact Affected versions: 12.1.0 or later
SYS-307854	**Fixed**: In rare cases, after performing hot upgrade while using Fibre Channel (FC) connectivity, the FC driver process might crash and then restart automatically. In some system or host configurations, this might result in hosts losing access to the storage system. Severity: High Impact Affected versions: 12.0.0 or later
SYS-307715	**Fixed**: In rare cases, when hosts are connected to the storage system via iSCSI, a Send Targets request (sent from a host) might cause a loss of access for all hosts connected to the storage system. Severity: High Impact Affected versions: 12.2.0 or later
SYS-307743	**Fixed**: In rare cases, after the recovery of a failed software component, the connectivity to a remote mirroring target might not be restored. Severity: High Impact Affected versions: 12.1.0 or later
SYS-307987	**Fixed**: In rare cases, when using asynchronous mirroring, specific write sequences (writes of 0) combined with a mirror deactivation procedure may result in loss of access upon re-activation. On even rarer occasions, loss of data can occur. Severity: High Impact Affected versions: 12.0.0 or later
SYS-308098	**Fixed**: In rare cases, if the hot upgrade of a storage system that is using encryption is canceled, an internal component failure might trigger an emergency shutdown procedure, following by loss of access by hosts to the storage system. Severity: High Impact Affected versions: 12.0.0 or later
SYS-304905	**Fixed**: When performing a system wipe-out operation and then disabling the system's data encryption (for example, for the purpose of moving the system off-premise), the system performs several emergency shutdowns before the encryption disabling is completed. Severity: Moderate Affected versions: 12.2.1 or later
SYS-306078	**Fixed**: In rare cases, I/O operations might be rejected due to an IOPS or bandwidth limit of a QoS performance class, even though these I/O operations are not subject to any QoS performance class limitation. Severity: Moderate Affected versions: 12.0.0 or later
SYS-306353	**Fixed**: Mirrored volumes that are not part of a consistency group (CG) cannot benefit from the 12.2.0 software enhancements that allow configuring an RPO of 1 minute and mirroring interval of 20 seconds. Severity: Moderate Affected versions: 12.2.0 or later
SYS-306627	**Fixed**: In rare cases, after a hot upgrade of a Primary system in a HyperSwap relationship over Fibre Channel (FC) connectivity, the FC driver process might crash and then restart automatically, causing temporary HyperSwap connectivity disconnection. Severity: Moderate Affected versions: 12.0.0 or later

Ticket ID	Description
SYS-307088	**Fixed**: In rare cases, the size of the volume in the data migration source might be incorrectly reported when using the **dm_define** or the **dm_activate** CLI commands. If this occurs, data migration cannot be performed. Severity: Moderate Affected versions: 12.0.3.b or later
SYS-307675	**Fixed**: In rare cases, internal disconnections between a grid controller and the flash enclosure might increase the latency. If this occurs, the `IO_TOOK_TOO_LONG` event is logged. Severity: Moderate Affected versions: 12.2.1 or later
SYS-307815	**Fixed**: In rare cases, a particular sequence of read and write I/O operations might cause performance degradation. Severity: Moderate Affected versions: 12.2.0 or later
SYS-280653	**Fixed**: Some CLI commands might hang when storage pools are being cleaned from snapshots. Severity: Service Affected versions: 12.0.0 or later
SYS-305312	**Fixed**: When using HyperSwap, the `HA_MASTER_UNAVAILABLE` and `HA_MASTER_AVAILABLE` system events are not logged for all volumes and consistency groups after a system reboot followed by disconnection from the peer system. As a result, incorrect values are shown when using the **ha_list** command as long as the disconnection continues. Severity: Service Affected versions: 12.1.0 or later
SYS-305901	**Fixed**: In rare cases, after a reboot or software upgrade of the storage system, the output of the **fan_list** CLI command might list fans as failed, even though the fans are fully operational. Severity: Service Affected versions: 12.0.0 or later
SYS-307419	**Fixed**: If the storage system stores secondary HyperSwap volumes, hot upgrade of this storage system might fail if the HyperSwap relation is not disabled beforehand. Severity: Service Affected versions: 12.2.1 or later
SYS-306886	**Fixed**: In the output of the **ipinterface_list_ports** CLI command, all iSCSI ports might be listed with the Link Up status, even though some ports have no connectivity. Severity: Low Affected versions: 12.0.0 or later

Fixes and features of earlier versions

IBM FlashSystem® A9000R 12.3.2.e includes all the features and fixes that were provided in previous 12.0.x, 12.1.x, 12.2.x, 12.3.0.x, 12.3.1.x, 12.3.2, 12.3.2.a, 12.3.2.a-1, 12.3.2.a-2, 12.3.2.a-3, 12.3.2.b, 12.3.2.c, and 12.3.2.e versions.

Limitations

As opposed to known issues, limitations are functionality restrictions that are part of the predefined system design and capabilities in a particular version.

Synchronous mirroring limitations

Use of synchronous remote mirroring is subject to the following:

- The reported size of data written on the master volume and slave volume are not equal, since zero data writes are not counted for the slave volume. This is by design and no actual data is lost.

Limitations with IBM SAN Volume Controller

When using FlashSystem A9000R as the back-end storage of IBM® SAN Volume Controller, the FlashSystem A9000R system is not aware of user data that is deleted or migrated on the front-end SAN Volume Controller. In addition, the FlashSystem A9000R physical capacity usage is not monitored by the front-end SAN Volume Controller.

As a result, the FlashSystem A9000R system cannot regularly reclaim freed capacity, and over time it may unexpectedly for the storage administrator run out of physical storage space.

To mitigate this limitation, when using FlashSystem A9000R as a back-end storage of an SAN Volume Controller, administrators are strongly advised to:

- Closely monitor the FlashSystem A9000R physical capacity using Hyper-Scale Manager or IBM Extended Command-Line Interface (XCLI) Utility.
- Keep track of FlashSystem A9000R capacity-related events.
- Upgrade to software version 8.1 or later of SAN Volume Controller.

For more information about SAN Volume Controller, see the SAN Volume Controller documentation.

Known issues

This section details the known issues in IBM FlashSystem® A9000R 12.3.2.e, along with possible solutions or workarounds (if available).

The following severity levels apply to known issues:

- **HIPER** – High Impact Pervasive. A critical issue that IBM® has either fixed or plans to fix promptly. Requires immediate customer attention or code upgrade.
- **High Impact** – Potentially irrecoverable error that might impact data or access to data in rare cases or specific situations/configurations.
- **Moderate** – Limited functionality issue and/or performance issue with a noticeable effect.
- **Service** – Non-disruptive recoverable error that can be resolved through a workaround.
- **Low** – Low-impact usability issue.

Important:

- **The issues listed below apply to version 12.3.2.e or earlier versions**. As long as a newer version has not yet been released, a newer release notes edition for version 12.3.2.e might be issued to provide a more updated list of known issues and workarounds.
- When a newer version is released for general availability, the release notes of version 12.3.2.e will no longer be updated. Accordingly, check the release notes of the newer version to learn whether any newly discovered issues affect version 12.3.2.e or whether the newer version resolves any of the issues listed below.

Ticket ID	Severity	Affected versions	Description
SYS-305091	High Impact	12.0.0 or later	In rare cases, using the **io_pause** command might cause loss of access to the storage system. Note: Starting from version 12.2.1, this CLI command is disabled until a fix for this issue becomes available. **Workaround**: No workaround is currently available.
SYS-305121	High Impact	12.1.0 or later	When using IBM Hyper-Scale Mobility and the volume migration is in Proxy state, a broken link to the destination system might cause hosts to lose access to one or more of the Proxy-state interface modules. **Workaround**: No workaround is currently available.
SYS-305998	High Impact	12.1.0 or later	When the storage system enters an out-of-physical-space state and the **SYSTEM_OUT_OF_PHYSICAL_SPACE** event is logged, freeing up physical space will allow the storage system to return to normal operation. However, in rare cases, the storage system might enter the same state again, even if there is some free space available. **Workaround**: Contact IBM Support.
SYS-305999	High Impact	12.0.0 or later	When an encrypted storage system is powered up, it may enter the maintenance state due to an issue in the encryption key management mechanism. **Workaround**: Contact IBM Support.
SYS-312000	High Impact	12.0.0 or later	In rare cases, when using the **mirror_delete** CLI command together with the force_on_slave=yes parameter on an asynchronous mirroring relation or a Multisite HA DR relation, while the volume is part of an active consistency group (CG), any asynchronous mirror volumes will not be updated, and miscellaneous management operations might fail and need to be retried. **Workaround**: Contact IBM Support.
SYS-282831	Moderate	12.0.2 or later	Performance improvements that were introduced in version 12.0.2 cannot be applied on data that was written to the system prior to upgrading to version 12.0.2. **Workaround**: When experiencing performance issues with data that was written to the system prior to upgrading to version 12.0.2, contact IBM Support.
SYS-285786	Moderate	12.1.0 or later	When unmapping a migrated volume (with IBM Hyper-Scale Mobility) or a HyperSwap® volume from a cluster, a delay of the host I/O response might be experienced, ranging from 20 to 170 milliseconds. **Workaround**: No workaround is currently available.
SYS-304840	Moderate	12.1.0 or later	Using the IBM Hyper-Scale Mobility (volume migration) feature between XIV® Gen3 systems and FlashSystem A9000 or A9000R systems is not supported during the hot upgrade of FlashSystem A9000 or A9000R. **Workaround**: No workaround is currently available.
SYS-304905	Moderate	12.2.0 or later	When performing a system wipe-out operation and then disabling the system's data encryption (for example, for the purpose of moving the system), the system performs several emergency shutdowns before the encryption disabling is completed. **Workaround**: Manually restart the system and allow the decryption process to complete.
SYS-309714	Moderate	12.1.0 or later	In a HyperSwap configuration, attaching a Quorum Witness to a domain is not possible. **Workaround**: No workaround is currently available.
SYS-314668	Moderate	12.3.2 or later	IP filtering is not enforced by the **IP_ACCESS** CLI command when the management ports are assigned with a VLAN ID. **Workaround**: No workaround is currently available.
SYS-314708	Moderate	12.0.0 or later	When using the **host_delete** CLI command on one of the secondary system ports designated for a HyperSwap volume only the secondary system port is removed and the port on the primary system remains configured. This can expose the volume to SCSI reservation conflict, causing unexpected access restrictions until the port is removed. **Workaround**: To avoid this issue, be sure to remove the port from the primary system before deleting the host, using the **host_remove_port** CLI command. If the port was not removed prior to using the **host_delete** CLI command, contact IBM support.

Ticket ID	Severity	Affected versions	Description
SYS-315170	Moderate	12.0.1 or later	In some cases when using a KeySecure key server, the **encrypt_disable** CLI command fails to complete and the `ENCRYPT_UNABLE_TO_UPDATE_KEY_DURING_DEACTIVATE_ON_KEYSERVER` event is emitted. **Workaround**: If the key already exists, and the `ENCRYPT_UNABLE_TO_UPDATE_KEY_DURING_DEACTIVATE_ON_KEYSERVER` is emitted, go to the Life Cycle properties dialog box of the key on the key server, and change the state from Pre-Active to Active. Once the key state is Active, repeat the **encrypt_disable** CLI command.
SYS-316178	Moderate	12.1.0 or later	In rare cases, a Fibre Channel zone change on the switch side may not be reflected accordingly on the FlashSystem A9000R side, due to a Fibre Channel driver issue in target mode. **Workaround**: Contact IBM Support.
SYS-279815	Service	12.0.0 or later	If a MicroLatency module is pulled out of the flash enclosure when the enclosure's internal redundancy state is either 'Synching' or 'Degraded', the management GUI indicates that the card has a spare, even if no spare is available. **Workaround**: Use the **flash_card_list** CLI command to display the correct status of spares.
SYS-280653	Service	12.0.0 or later	Some CLI commands might hang when storage pools are being cleaned from snapshots. **Workaround**: In such a case, wait until the command is carried out.
SYS-282223	Service	12.0.0 or later	Data migration over Fibre Channel (FC) does not work if the number of LUNs per host exceeds 255. **Workaround**: To resolve this issue, use 255 or less LUNs per host.
SYS-284723	Service	12.0.0 or later	Enabling encryption of a flash enclosure that contains a failed flash canister could lead to partial encryption of the flash enclosure, or to complete encryption failure. **Workaround**: To avoid this issue, before enabling encryption, verify that no flash enclosure contains a failed canister. Otherwise, contact IBM Support.
SYS-284298	Service	12.0.0 or later	When setting up LDAP authentication, entering the storage system password is required, but this requirement is not enforced. This could lead to LDAP authentication failure if the password is not provided. **Workaround**: To avoid this issue, always make sure to provide the storage system password in order to allow access to the storage system.
SYS-284512	Service	12.1.0 or later	An attempt to deactivate a Quorum Witness might render the Quorum Witness unresponsive, and remain in the *Deactivating* state without timing out (see also **SYS-285120** below). **Workaround**: To resolve this issue, proceed as follows in IBM Hyper-Scale Manager: 1. Remove the current Quorum Witness configuration. This action disassociates the Quorum Witness from the connected system(s). 2. Define a new Quorum Witness and make sure that the correct port is used. 3. Activate the newly defined Quorum Witness. 4. Connect the system(s), that were previously using the removed Quorum Witness, to the newly defined Quorum Witness.
SYS-285120	Service	12.1.0 or later	When connecting a system to a Quorum Witness, the Quorum Witness might become unresponsive, and remain in the *Activating* state without timing out. **Workaround**: To resolve this issue, proceed as follows in IBM Hyper-Scale Manager: 1. Deactivate the Quorum Witness. Note that the Quorum Witness state changes to *Deactivating* and remains *Deactivating* without timing out (see also **SYS-284512** above). 2. Remove the current Quorum Witness configuration. This action disassociates the Quorum Witness from the system(s). 3. Define a new Quorum Witness and make sure that the correct port is used. 4. Activate the newly defined Quorum Witness. 5. Connect the system(s), that were previously using the removed Quorum Witness, to the newly defined Quorum Witness.
SYS-286013	Service	12.0.0 or later	When an IBM Support technician is removing a canister from the flash enclosure for maintenance purposes, multiple events of `POD_IB_LINK_DETECTION_LINK_PERSISTENTLY_DISCONNECTED` might be unnecessarily emitted. These events can be safely ignored.
SYS-286449	Service	12.0.0 or later	When the CLI command **reservation_key_list** is applied to the entire system, the command output displays the maximum of 512 keys, even if the factual number of keys is greater. However, when applied to a volume, **reservation_key_list** outputs the entire key list. **Workaround**: Apply the command **reservation_key_list** per volume.
SYS-287068	Service	12.1.0 or later	In rare cases, the **ha_create** command might fail. **Workaround**: Retry the operation.
SYS-287305	Service	12.1.0 or later	If there are more than 200 volumes in a HyperSwap relation and a failover takes place, some of the `HA_AUTOMATIC_FAILOVER_SUCCESFUL` events might be lost. This does not affect the actual failover. **Workaround**: No workaround is currently available.
SYS-287279	Service	12.1.0 or later	When removing a target port and adding it back within less than 3 seconds, the high availability state for HyperSwap volumes is insufficient, and a `TARGET_CONNECTIVITY_HA_INSUFFICIENT` event is issued. **Workaround**: No workaround is currently available.
SYS-287533	Service	12.1.0 or later	Previously listed as SYS-310649. In rare cases, when adding a volume to a consistency group (CG) with the **cg_add_vol** command, the operation might fail with the event `MIRROR_LAST_SYNC_TIMES_DIFFER` recorded. **Workaround**: To avoid this issue, retry the **cg_add_vol** command after the mirroring synchronization is completed.

Ticket ID	Severity	Affected versions	Description
SYS-305312	Service	12.1.0 or later	When using HyperSwap, the `HA_MASTER_UNAVAILABLE` and `HA_MASTER_AVAILABLE` system events are not logged for all volumes and consistency groups after a system reboot followed by disconnection from the peer system. As a result, incorrect values are shown when using the **ha_list** command as long as the disconnection continues. **Workaround**: Wait until the high availability connection to the peer system is reestablished. The correct information will then be displayed.
SYS-305433	Service	12.1.0 or later	When using IBM Hyper-Scale Mobility and aborting (with **olvm_abort**) a migration process that is currently in Proxy state, the remote volume remains in locked read-only state. **Workaround**: Use the **vol_unlock** command to unlock the volume state.
SYS-305450	Service	12.0.0 or later	Manually created snapshots that have a reserved name prefix prevent the automatic creation of additional snapshots. **Workaround**: Delete and do not manually create snapshots that have a reserved name prefix.
SYS-305581	Service	12.0.0 or later	The **switch_mgmt_ip_list** CLI command does not work when using the additional switch argument for displaying the management IP address of a specific InfiniBand switch. **Workaround**: Use the **switch_mgmt_ip_list** command without the switch argument. This will display the management IP addresses of all InfiniBand switches.
SYS-309859	Service	12.1.0 or later	In some cases of desynchronization in the phase-out process of a grid element, the `node_failed` and `taken_over` events could be recorded. This does not impact any critical system operation. **Workaround**: No workaround is currently available.
SYS-310909	Service	12.2.1 or later	In rare cases, the Fibre Channel (FC) driver might crash and then immediately be restarted automatically. In such a case, the `FC_LINK_IS_NOW_UP` event is logged. **Workaround**: No workaround is currently available.
SYS-312467	Service	12.0.0 or later	In rare cases, the `TARGET_CONNECTION_HA_INSUFFICIENT` event appears, followed by the `TARGET_CONNECTION_HA_SUFFICIENT` event after a few seconds. This behavior has no actual impact. **Workaround**: No workaround is currently available.
SYS-313056	Service	12.1.0 or later	When using HyperSwap and a volume fails to be synchronized, the `VOLUME_IS_NOT_CONSISTENT_OLVM_DESTINATION` event is logged instead of the VOLUME_IS_NOT_CONSISTENT_HA_DESTINATION event. **Workaround**: No workaround is currently available.
SYS-276870	Low	12.0.0 or later	When a grid controller is replaced, or whenever the system is undergoing a software initiation process, data reduction saving statistics could be temporarily inaccurate. The statistics accuracy is regained as data is being accessed and processed again. This issue has a noticeable effect only when 1 TB or less has been written to the system. **Workaround**: No workaround is currently available.
SYS-306320	Low	12.0.0 or later	In rare cases, the process of enabling encryption may not complete successfully due to a timeout. If this occurs, the system emits the `ENCRYPT_ENABLE_NOT_COMPLETED` event with all the relevant details. **Workaround**: Retry activating encryption.
SYS-306814	Low	12.0.1 or later	By design, an attempt to set up IBM Hyper-Scale Mobility, when a volume with the same name as the source volume already exists at the destination system, fails due to an illegal volume name. However, the emitted return code (`REMOTE_VOLUME_NOT_APPLICABLE_FOR_OLVM`) provides a wrong reason for the failure. **Workaround**: No workaround is currently available.
SYS-306886	Low	12.0.0 or later	In the output of the **ipinterface_list_ports** CLI command, all iSCSI ports might be listed with the Link Up status, even though some ports have no connectivity. **Workaround**: No workaround is currently available.
SYS-313010	Low	12.0.0 or later	When iSCSI connectivity is used with a host, the `HOST_NO_MULTIPATH_ONLY_ONE_PORT` event might be logged, even though the host is connected through more than one iSCSI path. **Workaround**: No workaround is currently available.
SYS-314527	Low	12.3.0 or later	When a Multi-site HA DR relationship relation is established, it takes a few seconds for the Standby connectivity status to be updated. During this short period, no volume can be added to a consistency group. **Workaround**: Repeat the CLI command or GUI operation.
SYS-315110	Low	12.0.0 or later	The **ipinterface_run_arp** and **ipinterface_list_ips** CLI commands do not support iSCSI (IP-based) interface queries. **Workaround**: Use the **ipinterface_list** CLI command for iSCSI interface queries.
SYS-316921	Low	12.0.0 or later	Job size output for **mirror_statistics_get** incorrectly reports job sizes over 2 TiB. This behavior has no actual impact on system performance and the sync job can complete successfully. **Workaround**: No workaround is currently available.

Known interoperability issues

This section describes known issues or limitations that could arise when using IBM FlashSystem® A9000R in conjunction with other components in the production environment, such as hardware switches, host bus adapters (HBAs), other IBM® products, or different software platforms such as operating systems, clustering solutions, virtual machine servers, multipathing solutions, and other third-party components.

Ticket ID	Severity	Affected versions	Description

Ticket ID	Severity	Affected versions	Description
SYS-286719	Service	12.1.0 or later	Failover of a Windows Server 2008 or 2008 R2 cluster group results in a stop error on the new active node. This could occur when using these Windows Server versions with HyperSwap® volumes. No hotfix is currently available. **Workaround**: Refer to IBM Support guidelines for Windows Server 2008 and 2008 R2 as described in IBM Technote #S1010299.
SYS-286992	Service	12.1.0 or later	The HyperSwap feature is not compatible with VMware ESXi 6.0 Update 2 or earlier. **Workaround**: Use ESXi 6.0 Update 3 or later.
SYS-287993	Service	12.1.0 or later	Veritas Dynamic Multi-Pathing (DMP) in any configuration does not support HyperSwap volumes. **Workaround**: No workaround is currently available.

Release notes for IBM HyperSwap Quorum Witness

The following release information is available for IBM Spectrum Accelerate Family HyperSwap® Quorum Witness 1.0.3.

What's new in 1.0.3

IBM Spectrum Accelerate Family HyperSwap® Quorum Witness 1.0.3 introduces new Linux® version support.

Version: 1.0.3

General availability date: 2 February 2021

New support for Red Hat® Enterprise Linux (RHEL) and CentOS Linux 8.2 and 8.4

Quorum Witness 1.0.3 now supports RHEL and CentOS versions 8.2 and 8.4.
Important: If your operating system is lower than 8.2, either upgrade your RHEL or CentOS version accordingly, or use Quorum Witness 1.0.2. The operating system upgrade must be done prior to Quorum Witness 1.0.3 installation. See the Upgrading an existing installation section of the User Guide for more information.

Compatibility and requirements

The IBM Spectrum Accelerate Family HyperSwap® Quorum Witness is compatible with different versions of the Linux® operating system and the storage systems.

Certain software packages, patches, or drivers must be installed on the host, as detailed in the following subsections.

- **Supported Linux versions**
  IBM Spectrum Accelerate Family HyperSwap Quorum Witness 1.0.3 supports different Linux versions and editions, as listed in the following table.
- **Supported storage systems**
  IBM Spectrum Accelerate Family HyperSwap Quorum Witness 1.0.3 supports IBM FlashSystem® A9000 and A9000R with microcode version 12.1.x, 12.2.x, and 12.3.x.
- **Host requirements**
  The IBM Spectrum Accelerate Family HyperSwap Quorum Witness requires a physical or virtual (guest) server host. The host must meet the following minimum requirements.

Supported Linux versions

IBM Spectrum Accelerate Family HyperSwap® Quorum Witness 1.0.3 supports different Linux® versions and editions, as listed in the following table.

Operating system	Architecture
Red Hat® Enterprise Linux (RHEL) 8.2, 8.4	x86_64
CentOS Linux 8.2, 8.4	x86_64

Supported storage systems

IBM Spectrum Accelerate Family HyperSwap® Quorum Witness 1.0.3 supports IBM FlashSystem® A9000 and A9000R with microcode version 12.1.x, 12.2.x, and 12.3.x.

Host requirements

The IBM Spectrum Accelerate Family HyperSwap® Quorum Witness requires a physical or virtual (guest) server host. The host must meet the following minimum requirements.

Important: Quorum Witness cannot reside on the same set of FlashSystem A9000 and/or A9000R systems of which it serves as part of a HyperSwap configuration.

- 64-bit dual-core CPU, minimum 2.3 GHz
- 4 GB of RAM
- Hard or solid-state drive:
  - SAS/SATA interface
  - 7200 RPM speed, minimum 100 MB/s (hard disk drive only)
  - RAID protection
  - 40 GB of free space
- End-to-end network connectivity:
  - Connectivity between all grid controller management ports and the Quorum Witness
  - Bandwidth – 6 Mbps multiplied by the number of registered storage systems
  - Latency – 0.75 seconds
  - Packet loss – 0.1%

Note:

- For best performance, 99.999% (five-nines) network availability is recommended for all elements of HyperSwap solution. This includes the storage systems and the Quorum Witness host.
- To decrease packet loss and latency on the network, use assured forwarding (AF) packet tagging.
- To achieve high availability of the Quorum Witness, it must be installed on a fault-tolerant virtual machine. Refer to the relevant VMware documentation for details on vSphere Fault Tolerance (FT) solution.
- If a VM is configured, be sure to reserve the relevant resources, as specified in this section.

Change log

This section summarizes the changes made in different version releases of the IBM Spectrum Accelerate Family HyperSwap® Quorum Witness.

Note: New functional features of these versions are only briefly summarized. For a detailed summary of the new functional features of a specific version, refer to the 'What's new' section of the release notes.

- **1.0.3 (February 2021)**
  IBM Spectrum Accelerate Family HyperSwap Quorum Witness 1.0.3 replaces version 1.0.2, providing new Linux® support and resolves the following issue.
- **1.0.2 (August 2019)**
  Version 1.0.2 replaced version 1.0.1, improving resiliency and Reliability Availability and Serviceability (RAS), and resolved the following issues.
- **1.0.1 (June 2018)**
  Version 1.0.1 replaced version 1.0.0, improving the installation procedure of the software package, as detailed below.
- **1.0.0 (June 2017)**
  Version 1.0.0 of the IBM Spectrum® Accelerate Family HyperSwap Quorum Witness was the first released version, intended for use IBM FlashSystem® A9000 and A9000R storage systems version 12.1 deployed in the HyperSwap applications.

1.0.3 (February 2021)

IBM Spectrum Accelerate Family HyperSwap® Quorum Witness 1.0.3 replaces version 1.0.2, providing new Linux® support and resolves the following issue.

Ticket ID	Description
QUOR-262	**Fixed**: When the CVE-2018-16888 patch is installed on a Red Hat® Enterprise Linux (RHEL) or CentOS Linux server, some processes are unable to run, and the Quorum Witness is unable to work properly.

1.0.2 (August 2019)

Version 1.0.2 replaced version 1.0.1, improving resiliency and Reliability Availability and Serviceability (RAS), and resolved the following issues.

Ticket ID	Description
QUOR-258	**Fixed**: In some cases, an upgrade to the Linux® kernel or cloning/migrating the Virtual Machine may cause a change in the system_uid of the system, preventing Quorum Witness from starting.

1.0.1 (June 2018)

Version 1.0.1 replaced version 1.0.0, improving the installation procedure of the software package, as detailed below.

Ticket ID	Description
QUOR-199	**Enhancement**: The ibmqw user home directory can be customized prior to the installation of the Quorum Witness software.

1.0.0 (June 2017)

Version 1.0.0 of the IBM Spectrum® Accelerate Family HyperSwap® Quorum Witness was the first released version, intended for use IBM FlashSystem® A9000 and A9000R storage systems version 12.1 deployed in the HyperSwap applications.

Limitations

As opposed to known issues, limitations are functionality restrictions that are part of the predefined system design and capabilities in a particular version.

Important: Quorum Witness cannot reside on the same set of FlashSystem A9000 and/or A9000R systems of which it serves as part of a HyperSwap® configuration.

Operational capacity

Quorum Witness-related capacity of the IBM® HyperSwap solution:

- Maximum number of registered storage systems per Quorum Witness – 12.
- Maximum number of Quorum Witnesses defined on a storage system – 2 (one active and one standby).

Check for newer release note publications to learn about limitations that could be discovered and advertised in the future.

Known issues

This section details the known issues in IBM Spectrum Accelerate Family HyperSwap® Quorum Witness 1.0.3, along with possible solutions or workarounds (if available).

The following severity levels apply to known issues:

- **HIPER** – High Impact Pervasive. A critical issue that IBM® has either fixed or plans to fix promptly. Requires immediate customer attention or code upgrade.
- **High Impact** – Potentially irrecoverable error that might impact data or access to data in rare cases or specific situations/configurations.
- **Moderate** – Limited functionality issue and/or performance issue with a noticeable effect.
- **Service** – Non-disruptive recoverable error that can be resolved through a workaround.
- **Low** – Low-impact usability-related issue.

Important:

- **The issues listed below apply to version 1.0.3 or earlier versions**. As long as a newer version has not yet been released, a newer release notes edition for version 1.0.3 might be issued to provide a more updated list of known issues and workarounds.
- When a newer version is released for general availability, the release notes of version 1.0.3 will no longer be updated. Accordingly, check the release notes of the newer version to learn whether any newly discovered issues affect version 1.0.3 or whether the newer version resolves any of the issues listed below.

Ticket ID	Severity	Affected versions	Description
QUOR-248	High Impact	1.0.0 or later	Mounting /tmp with *noexec* option results in package installation failure. In the result of failure, the message `License not accepted` is displayed instead of the license being displayed. **Workaround:** To resolve this issue, mount /tmp with *exec* option, then remount it with *noexec* after successful installation of the IBM Spectrum Accelerate Family HyperSwap Quorum Witness.
QUOR-268	High Impact	1.0.0 or later	In extremely rare cases, when a Quorum Witness instance is connected to 8 or more FlashSystem A9000 or A9000R systems that have all experienced simultaneous multiple failover scenarios, the response time of the Quorum Witness may experience increased latency, resulting in data unavailability. **Workaround:** If simultaneous multiple failovers on 8 or more systems has occurred, contact IBM Support.

QUOR-260	Service	1.0.2 or later	When the Linux® operating system is setup and /etc is mounted on a different filesystem then /opt, creating hard links during Quorum Witness installation fails.
			Workaround: To avoid this issue, configure a single file system before installing Quorum Witness.
			
			If Quorum Witness has already been installed, stop the services and manually copy the following files:
			• cp -pf /opt/ibm/ibm_quorum_witness/conf.d/ibmsyslog.conf /etc/rsyslog.d/40-ibmsyslog.conf
			• cp -pf /opt/ibm/ibm_quorum_witness/conf.d/nginx/qw_nginx.conf /etc/nginx/conf.d/qw_nginx.conf
			
			After these files are copied, restart the services as follows:
			`systemctl restart nginx`
			`systemctl restart ibm_quorum_witness`

Release notes for IBM Hyper-Scale Manager

Use IBM® Hyper-Scale Manager to manage IBM Spectrum Accelerate systems.

For release information about the various IBM Hyper-Scale Manager version releases, see IBM Hyper-Scale Manager documentation.

Release notes for IBM XCLI Utility

The following release information is available for version 5.5.4.1 of IBM® Extended Command-line Interface (XCLI) Utility.

Overview

IBM® Extended Command-Line Interface (XCLI) Utility provides a terminal-based interface for issuing storage system management, monitoring, and maintenance commands from a client computer upon which the Utility is installed.

For detailed information about the command-line interface (CLI) commands for a particular storage system, see the CLI Reference Guide of that storage system.

The IBM XCLI Utility software package can be downloaded at any time from the IBM Fix Central website (www.ibm.com/support/fixcentral).

What's new in 5.5.4.1

IBM® XCLI Utility 5.5.4.1 is minor release that resolves a compatibility issue with Linux operating systems.

General availability date: June 2022

For more information about the resolved issue, see 5.5.4.1 (June 2022).

Compatibility

This section details the compatibility of IBM® XCLI Utility 5.5.4.1 with storage system versions and operating system versions.

Supported storage systems

IBM XCLI Utility 5.5.4.1 supports the following storage system versions:

Storage system	Microcode version
IBM FlashSystem® A9000R	12.3.x
IBM FlashSystem A9000	12.3.x
IBM XIV® Storage System Gen3 (models 281x-214 and 281x-314)	11.6.x
IBM Spectrum Accelerate	11.5.3.c, 11.5.4

Note:

- For the latest XIV Storage System code recommendations, see IBM XIV Storage Systems Code Update Recommendation (ibm.com®/support/pages/ibm-xiv-storage-systems-code-update-recommendation).
- For the latest FlashSystem A9000 and A9000R code recommendations, see IBM FlashSystem A9000/A9000R Code Update Recommendation (ibm.com/support/pages/node/707025).

Supported operating systems

IBM XCLI Utility 5.5.4.1 supports the following operating system versions:

Operating system	Bit version / CPU architecture
Red Hat® Enterprise Linux® (RHEL) 7.x, 8.x	64-bit
Microsoft Windows 10	32-bit or 64-bit

Note: For details on how to install the XCLI Utility, see the Command-Line Interface (CLI) Reference Guide.

Change log

This change log summarizes the changes that were made in recent 5.x code level versions of IBM® XCLI Utility.

Note: New functional features of these versions are only briefly summarized. For a detailed summary of the new functional features of a specific version, refer to the 'What's new' section of its individual release notes.

- **5.5.4.1 (June 2022)**
  IBM XCLI Utility 5.5.4.1 resolves the following issue.
- **5.5.4 (May 2022)**
  IBM XCLI Utility 5.5.4 introduced new Red Hat® Enterprise Linux® (RHEL) 7.x and 8.x support and resolved the following issue.

- **5.2.0.a (July 2017)**
  IBM XCLI Utility 5.2.0.a resolved the following compatibility issue.
- **5.2.0 (June 2017)**
  IBM XCLI Utility 5.2.0 updates the supported Java™ version to 7 update 131, and thereby resolves security issues in the previously supported Java version. In addition, IBM XCLI Utility 5.2.0 introduces an important usability enhancement, as listed in the following table.
- **5.1.0 (November 2016)**
  IBM XCLI Utility 5.1.0 introduced a number of security enhancements and resolves a few miscellaneous issues, as listed in the following table.

5.5.4.1 (June 2022)

IBM® XCLI Utility 5.5.4.1 resolves the following issue.

Ticket ID	Description
PUI-17822	**Fixed**: The IBM XCLI Utility logs for Linux systems have increased log verbosity.

5.5.4 (May 2022)

IBM® XCLI Utility 5.5.4 introduced new Red Hat® Enterprise Linux® (RHEL) 7.x and 8.x support and resolved the following issue.

Ticket ID	Description
PUI-17810	**Fixed**: Common security issue. For more information about this security issue, see CVE-2021-4104 on the Common Vulnerabilities and Exposures (CVE) information website (cve.org).

5.2.0.a (July 2017)

IBM® XCLI Utility 5.2.0.a resolved the following compatibility issue.

Ticket ID	Description
PUI-13473	**Fixed**: The IBM XCLI Utility installation on Windows 10 fails with an error message.

5.2.0 (June 2017)

IBM® XCLI Utility 5.2.0 updates the supported Java™ version to 7 update 131, and thereby resolves security issues in the previously supported Java version. In addition, IBM XCLI Utility 5.2.0 introduces an important usability enhancement, as listed in the following table.

Ticket ID	Description
PUI-12222	**Enhancement**: The font size in the XCLI Utility window is now customizable. To increase the font size, press CTRL+. To decrease the font size, press CTRL-.
PUI-12015	**Enhancement**: The supported Java version is updated to 7 update 131. This Java version resolves security issues in the previously supported Java version.

5.1.0 (November 2016)

IBM® XCLI Utility 5.1.0 introduced a number of security enhancements and resolves a few miscellaneous issues, as listed in the following table.

Ticket ID	Description
UI-265822	**Enhancement**: The direct command line login in the XCLI Utility can now be disabled. To disable it:, in the GUI12\xiv-general.properties file, do one of the following: • Change the value of the parameter xcli_direct_login_enabled to false • Add the string xcli_direct_login_enabled=false Note: Prior to modifying the xcli_direct_login_enabled parameter, make sure that the XIV® GUI is closed. Otherwise, the changes will not be applied. If the direct login is disabled, all direct authentication attempts will fail (command line arguments -u and -p are prohibited).
UI-265821	**Enhancement**: If the credentials file is created with the --protect argument, only the user running the command has read/write permissions to this file. No permissions are granted to any other user.
UI-265820	**Enhancement**: Starting from version 5.1, user passwords are not retained in XCLI or GUI log files.
UI-265819	**Enhancement**: Weak cryptographic algorithms used for obfuscation of user credentials in previous versions have been thoroughly enhanced and strengthened.
UI-265818	**Security enhancement**: User credentials are now required every time when launching the XCLI Utility or XIV Top from the XIV GUI.

Limitations

Currently there are no functional limitations in IBM® XCLI Utility 5.5.4.1.

Known issues

Currently there are no known issues in IBM® XCLI Utility 5.5.4.1.

Release notes for IBM Remote Support Proxy

The following release information is available for version 1.3.2.2 of IBM® Remote Support Proxy.

What's new in 1.3.2.2

IBM Remote Support Utility 1.3.2.2 is a minor release update, providing enhancements to the utility.

General availability date: May 2022

New front servers

This version adds new front servers. The following is the current list of IP addresses for all available front servers:

- 129.33.206.139
- 204.146.30.139
- 170.225.126.11 (new)
- 170.225.126.12 (new)
- 170.225.127.11 (new)
- 170.225.127.12 (new)

New Red Hat Enterprise Linux (RHEL) 8.x support

Remote Support Utility 1.3.2.2 introduces Red Hat Enterprise Linux (RHEL) 8.x operating system support for x86-64 architectures.

New SUSE Linux Enterprise Server (SLES) 15 support

Remote Support Utility 1.3.2.2 introduces SUSE Linux® Enterprise Server (SLES) 15 or later operating system support for x86-64 architectures.

Supported operating systems

IBM® Remote Support Proxy utility 1.3.2.2 supports different Linux® versions and editions, as listed in the following table.

Operating system	Architecture
Red Hat Enterprise Linux (RHEL) 7.x	x86-64
Red Hat Enterprise Linux (RHEL) 8.x	x86-64
SUSE Linux Enterprise Server (SLES) 12 or later	x86 and x86-64
SUSE Linux Enterprise Server (SLES) 15 or later	x86 and x86-64

For all the supported operating systems, the minimum required version of OpenSSL is openssl-1.0.2.

Supported storage systems

IBM® Remote Support Proxy utility 1.3.2.2 supports different microcode versions of the IBM storage systems, as listed in the following table.

Storage system	Microcode version
IBM XIV® Storage System	11.6.2.x
IBM Spectrum Accelerate	11.5.3.c, 11.5.4
IBM FlashSystem® A9000	12.3.2.x
IBM FlashSystem A9000R	12.3.2.x

Note:

- Newer microcode versions may also be compatible. For questions about compatibility, contact IBM Support.
- For the latest XIV Storage System code recommendations, see IBM XIV Storage Systems Code Update Recommendation (ibm.com/support/pages/ibm-xiv-storage-systems-code-update-recommendation).
- For the latest FlashSystem A9000 and A9000R code recommendations, see IBM FlashSystem A9000/A9000R Code Update Recommendation (ibm.com/support/pages/node/707025).

Limitations

At the time of publication, there were no known limitations.

Known issues

This section details the known issues in IBM Remote Support Proxy utility 1.3.2.2, along with possible solutions or workarounds (if available).

The following severity levels apply to known issues:

- **HIPER** – High Impact Pervasive. A critical issue that IBM® has either fixed or plans to fix promptly. Requires immediate customer attention or code upgrade.
- **High Impact** – Potentially irrecoverable error that might impact data or access to data in rare cases or specific situations/configurations.
- **Moderate** – Limited functionality issue and/or performance issue with a noticeable effect.
- **Service** – Non-disruptive recoverable error that can be resolved through a workaround.
- **Low** – Low-impact usability-related issue.

Important:

- **The issues listed below apply to version 1.3.2.2 or earlier versions**. As long as a newer version has not yet been released, a newer release notes edition for version 1.3.2.2 might be issued to provide a more updated list of known issues and workarounds.
- When a newer version is released for general availability, the release notes of version 1.3.2.2 will no longer be updated. Accordingly, check the release notes of the newer version to learn whether any newly discovered issues affect version 1.3.2.2 or whether the newer version resolves any of the issues listed below.

Ticket ID	Severity	Description
PXY-235410	High Impact	Mounting /tmp with *noexec* option results in package installation failure. In the result of failure, the message `License not accepted` is displayed instead of the license being displayed. **Workaround:** To resolve this issue, mount /tmp with *exec* option, then remount it with *noexec* after successful installation of the IBM Remote Support Proxy utility.
PXY-235418	Service	In some cases, during installation, the license is not displayed and the application terminates unexpectedly due to a missing bzip2 package on your system. **Workaround:** Install the bzip2 package using the Linux package manager and try to install the Remote Support Proxy again.

Release note PDF publications

Use this page to access PDF versions of release notes for IBM FlashSystem® A9000R, IBM Spectrum Accelerate Family HyperSwap Quorum Witness, IBM® Hyper-Scale Manager, IBM Extended Command-Line Interface (XCLI) Utility, and IBM Remote Support Proxy.

- IBM FlashSystem A9000R release notes
- IBM Spectrum Accelerate Family HyperSwap Quorum Witness release notes
- IBM Hyper-Scale Manager release notes
- IBM XCLI Utility release notes
- IBM Remote Support Proxy release notes

To view a PDF file, you need Adobe Reader. You can download it at no charge from the Adobe website.

IBM FlashSystem A9000R release notes

Release	Edition	PDF download link
12.3.2.e	First edition (January 2023)	Download
12.3.2.d	First edition (March 2022)	Download
12.3.2.c	Second Edition (March 2022)	Download
12.3.2.b	Third Edition (January 2021)	Download
12.3.2.a-3	Second Edition (January 2021)	Download
12.3.2.a-1	Sixth Edition (January 2021)	Download
12.3.2.a	Sixth Edition (January 2020)	Download
12.3.2	Fifth Edition (November 2019)	Download
12.3.1	Fourth Edition (November 2019)	Download
12.3.0.a	Sixth Edition (November 2019)	Download
12.3.0	Fifth Edition (November 2019)	Download

IBM Spectrum Accelerate Family HyperSwap® Quorum Witness release notes

Release	Edition	PDF download link
1.0.3	Third Edition (August 2021)	Download
1.0.2	Fourth Edition (February 2021)	Download
1.0.1	Second Edition (July 2019)	Download

IBM Hyper-Scale Manager release notes

To access PDF versions of release notes publications, see the Publications and related information within the IBM Hyper-Scale Manager documentation.

IBM XCLI Utility release notes

Release	Edition	PDF download link
5.5.4.1	First Edition (June 2022)	Download
5.5.4	First Edition (May 2022)	Download
5.2.0.a	Fourth Edition (May 2019) (Windows OS only)	Download
5.2.0	Eighth Edition (July 2020)	Download

IBM Remote Support Proxy release notes

Release	Edition	PDF download link
1.3.2.2	Second Edition (January 2023)	Download
1.3.2.1	Seventh Edition (July 2021)	Download

Product overview

The following overview information describes the IBM FlashSystem® A9000R functional features.

Introduction

IBM FlashSystem® A9000R is a high-end, grid scale, all-flash storage platform that delivers ultra-fast storage together with mission-critical features, including inline deduplication and data reduction by compression, smart scaling, distributed data, automatic load balancing, and a multitude of advanced enterprise-class features and capabilities.

Figure 1. IBM FlashSystem A9000R unit

As a pre-integrated rack-based offering, IBM FlashSystem A9000R aggregates grid elements (each containing two grid controllers and one flash enclosure) within a 42U integrated rack solution. Grid elements are interconnected by integrated InfiniBand switches, forming a scale-out grid fabric that delivers consistent, predictable high performance and ultra-low latency, even under heavy workloads with full data reduction enabled. The grid architecture enables the system to maintain this performance autonomously by evenly distributing every workload's data across all the system resources in real time. QoS features also help ensure that tenant service levels are not compromised.

In its core architecture, IBM FlashSystem A9000R utilizes IBM FlashCore® technology together with IBM® MicroLatency® modules, providing high density, low latency, and high IOPS performance. In addition, IBM FlashSystem Enhanced Endurance Technology reduces flash disk wear-out and ensures long-term durability of the flash storage components, even under heavy workloads.

On top of that, the built-in Spectrum Accelerate™ software provides redirect-on-write space-efficient snapshots, along with asynchronous and synchronous replication to enable granular data protection without increasing costs. To deliver high availability levels that exceed the IBM FlashSystem A9000R 99.999 availability, the system embeds a native implementation of IBM HyperSwap® capability, delivering active-active data access and transparent failover per volume, across IBM FlashSystem A9000R and IBM FlashSystem A9000 arrays and across data center.

Built as a rack-based grid scale system, IBM FlashSystem A9000R is able to meet your diverse workload requirements with a choice of performance and capacity optimized configurations. The IBM Hyper-Scale Manager allows orchestration of private and hybrid multi-tenant cloud environments at very large scales. The IBM Hyper-Scale Manager provides the ability to manage multiple IBM FlashSystem A9000R, IBM FlashSystem A9000, IBM XIV® and IBM Spectrum Accelerate solutions from a single pane of glass.

IBM FlashSystem A9000R is ideal for large enterprises that rely on fast, redundant, and high-capacity data storage, offering high service levels for dynamic workloads and easy hyper-scaling, while supporting multi-tenant environments, flexible consumption models, and robust cloud automation and integration capabilities. IBM FlashSystem A9000R offers security and data protection through advanced remote mirroring, hot encryption, and self-healing mechanisms.

Cloud economics and agility are adopted with ready A9000R solutions for Kubernetes and iSwarm container frameworks, as well as IBM Cloud Private, VMware, OpenStack, and Microsoft environments, at no additional cost.

- **Architecture**
  The high-end architecture of IBM FlashSystem A9000R comprises both the hardware architecture and logical architecture.
- **Functionality**
  IBM FlashSystem A9000R delivers a range of advanced storage functional features.
- **Specifications**
  For the IBM FlashSystem A9000R hardware and other specifications, refer to the marketing data sheet or to the IBM FlashSystem A9000R Deployment Guide.
- **Performance**
  IBM FlashSystem A9000R provides dynamic performance and ultra-low latency of IBM FlashCore technology, integrated with an innovative grid architecture, extensive scalability, and industry-leading IBM software-defined storage capabilities.

Architecture

IBM FlashSystem® A9000R comprises both the hardware architecture and logical architecture.

The primary system components are referred to as *grid elements*. This term describes a combination of one flash enclosure and two grid controllers within FlashSystem A9000R machines. With the exception of model 425 and U25 minimum configurations, all FlashSystem A9000R grid element valid configurations contain a number of flash enclosures and double that number of grid controllers.

- **Model 415**
  Model 415 rack configurations have a range between a minimum of two grid elements to a maximum of six grid elements.
- **Models 425 and U25**
  Model 425 and U25 rack configurations have a range between a minimum configuration of three grid controllers and one flash enclosure (one grid element with an additional grid controller) to a maximum of four grid elements.
- **Flash enclosures**
  Each flash enclosure comprises two fully redundant flash canisters, 12 MicroLatency® flash modules, two battery modules, and two power supply units.
- **Grid controllers**
  Each grid controller is a high-end Intel Xeon-based compute server, which acts as a core component that provides the interface and advanced compute functions.
- **Back-end interconnect**
  Internal communication (interconnect) between grid controllers or flash enclosures is carried over a 56 Gbps Fourteen Data Rate (FDR) InfiniBand network infrastructure, with full redundancy.
- **Logical architecture**
  The IBM FlashSystem A9000R logical architecture is built on the IBM Spectrum Accelerate software, with many added enhancements to optimize the software stack for use with flash storage.
- **Scale-out grid architecture**
  The IBM FlashSystem A9000R scale-out grid architecture enables the increase of the system's capacity and performance.

Model 415

Model 415 rack configurations have a range between a minimum of two grid elements to a maximum of six grid elements.

A minimum rack configuration is comprised of two grid elements and common components.

Figure 1. Model 415 minimum rack configuration

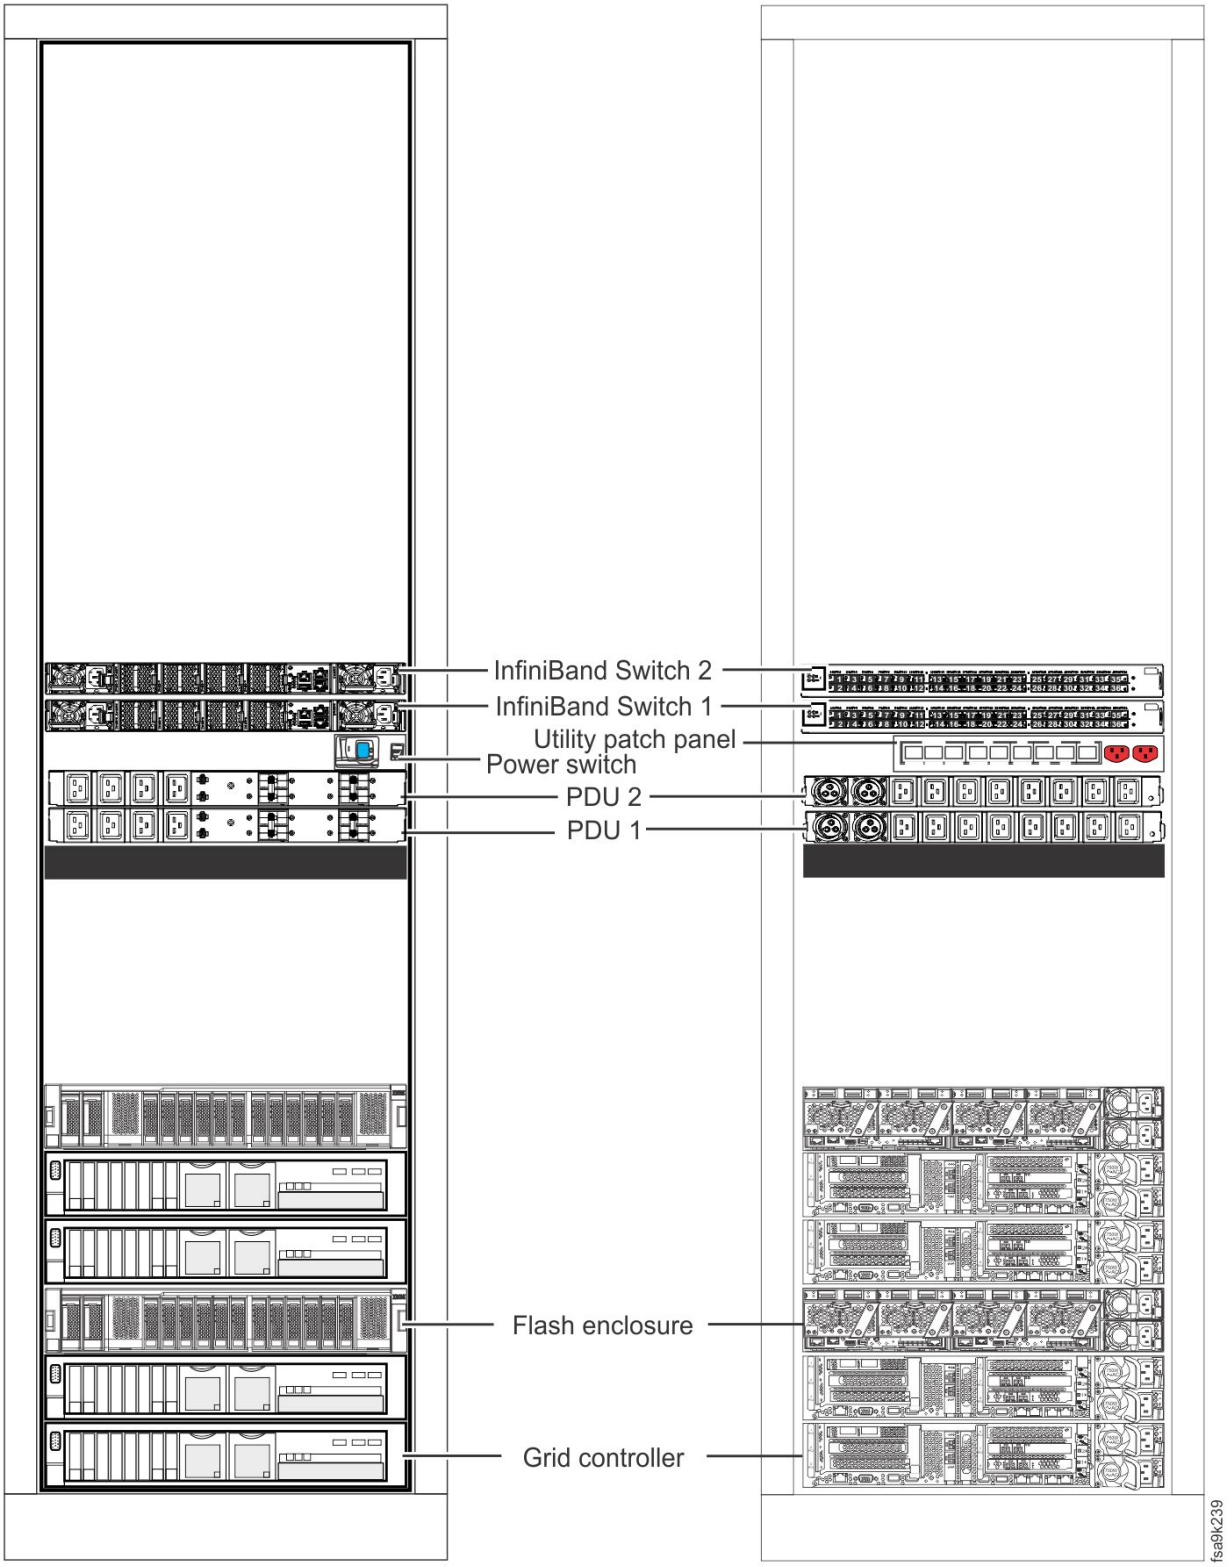

Partially populated rack configurations contain three, four, or five grid elements and common components. The maximum configuration (full rack) is comprised of six grid elements and common components.

Figure 2. Model 415 full rack configuration

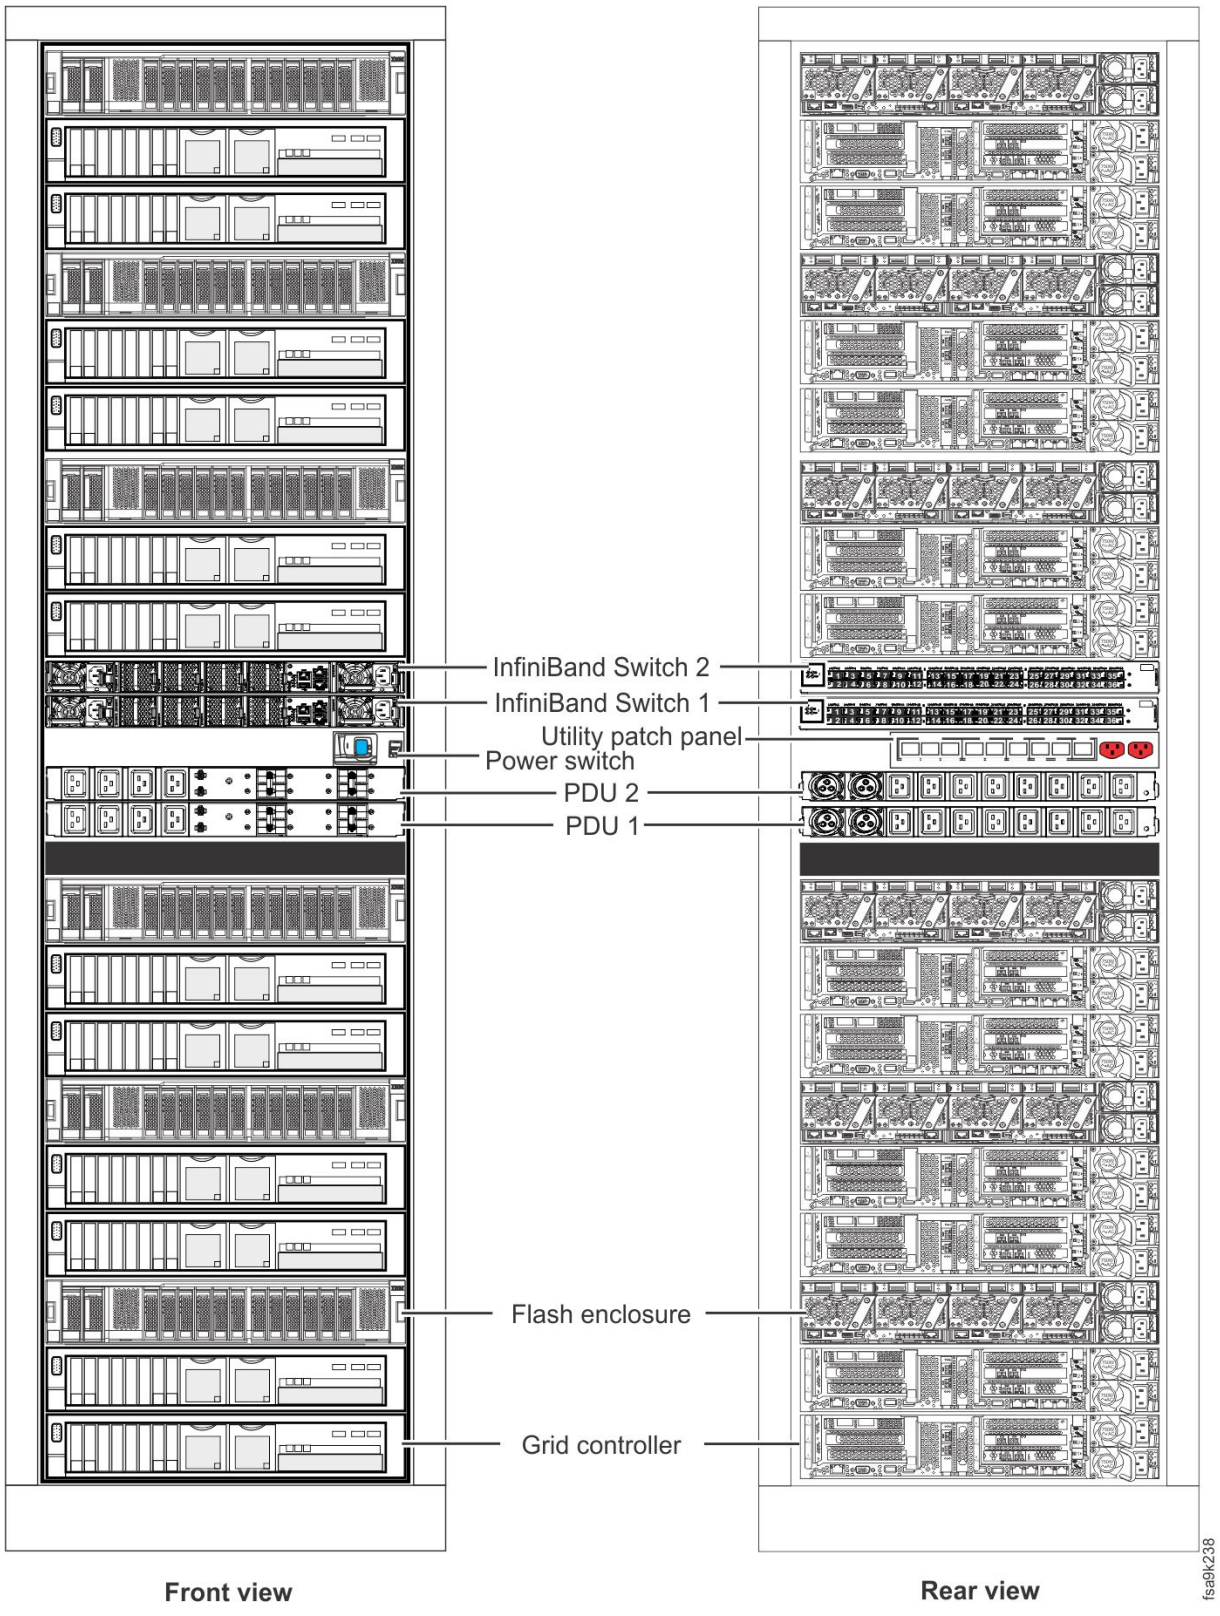

All flash enclosures contain 12 IBM MicroLatency modules. Depending on the flash enclosure ordered, MicroLatency modules are available with 2.9 TB or 5.7 TB raw capacity

Models 425 and U25

Model 425 and U25 rack configurations have a range between a minimum configuration of three grid controllers and one flash enclosure (one grid element with an additional grid controller) to a maximum of four grid elements.

Starting from software version 12.3.0, a minimum rack configuration is comprised of three grid controllers and one flash enclosure (one grid element and an additional grid controller) and common components. Until software version 12.3.0, a minimum rack configuration comprised of two grid elements, each containing three grid controllers and one flash enclosure.

Figure 1. Model 425 minimum rack configuration

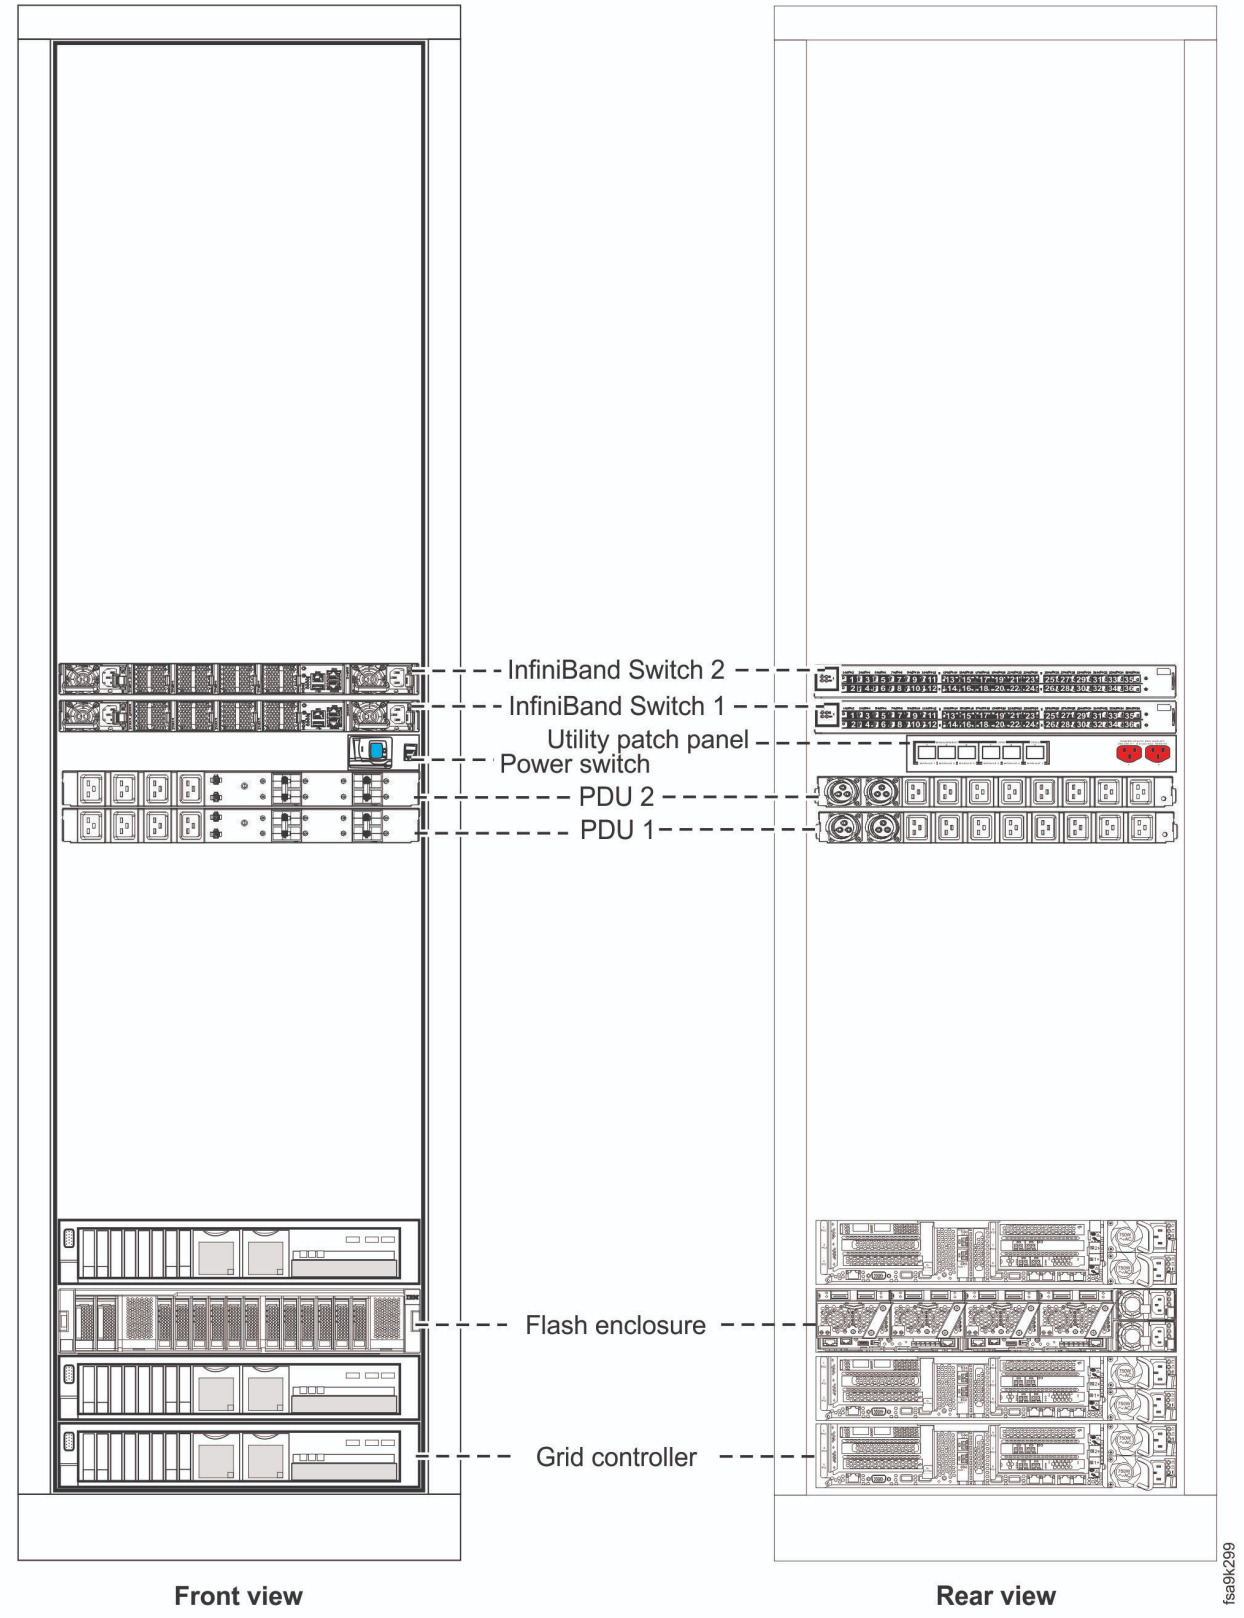

Starting from software version 12.3.1, non-disruptive scale-out from this minimum rack configuration is supported.

Partially populated rack configurations contain two or three grid elements and common components. The maximum configuration (full rack) is comprised of four grid elements and common components.

Figure 2. Models 425 and U25 full rack configuration

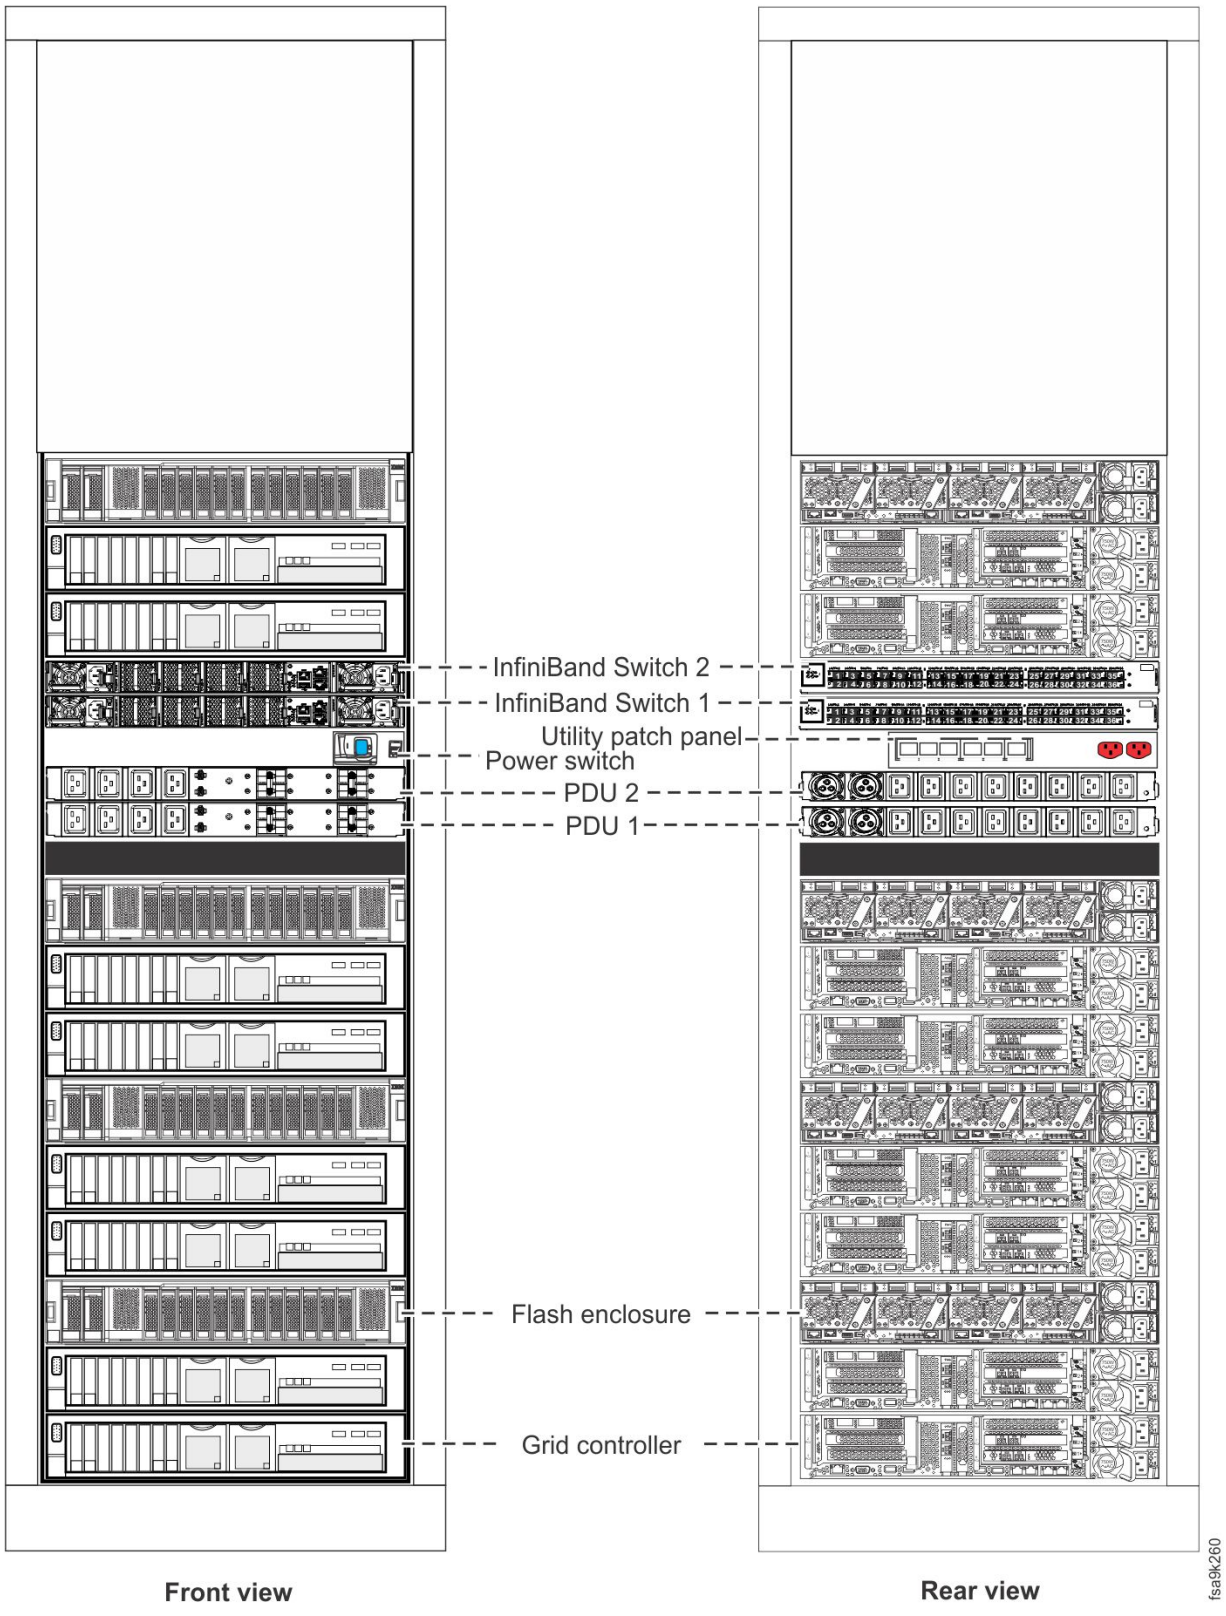

All flash enclosures contain 12 IBM MicroLatency modules. Depending on the flash enclosure ordered, MicroLatency modules are available with 3.6 TB, 8.5 TB, or 18 TB raw capacity.

Note: Data is automatically processed and distributed through all grid elements that are included in a FlashSystem A9000R system, and is not limited to any particular set of two grid controllers and one flash enclosure. Starting from software version 12.0.3, IBM FlashSystem® A9000R supports non-disruptive addition of sets of grid element sets that include two grid controllers and one flash enclosure, up to the full physical capacity of the rack system.

Flash enclosures

Each flash enclosure comprises two fully redundant flash canisters, 12 MicroLatency® flash modules, two battery modules, and two power supply units.

Each of the two flash canisters includes a RAID 5 controller, interface controller, management ports, and fans. The two redundant canisters are hot-swappable.

The usable capacity of MicroLatency flash modules is shown in the following table:

Table 1. Usable capacity of MicroLatency flash modules

Model 415	Models 425 and U25
12 hot-swap 1.2, 2.9, or 5.7 TB IBM® MicroLatency modules	12 hot-swap 3.6, 8.5, or 18 TB IBM MicroLatency modules

The maximum usable RAID 5 array capacity is 180 TB.

The flash enclosures of model 425 feature 3D TLC (Triple Level Cell) NAND MicroLatency modules, enabling dramatic improvement in density and TCO (total cost of ownership).

IBM FlashCore technology

The flash enclosures deliver the full range of the unique IBM FlashCore® technology, as summarized in the following figure.

Figure 1. IBM FlashCore technology benefits

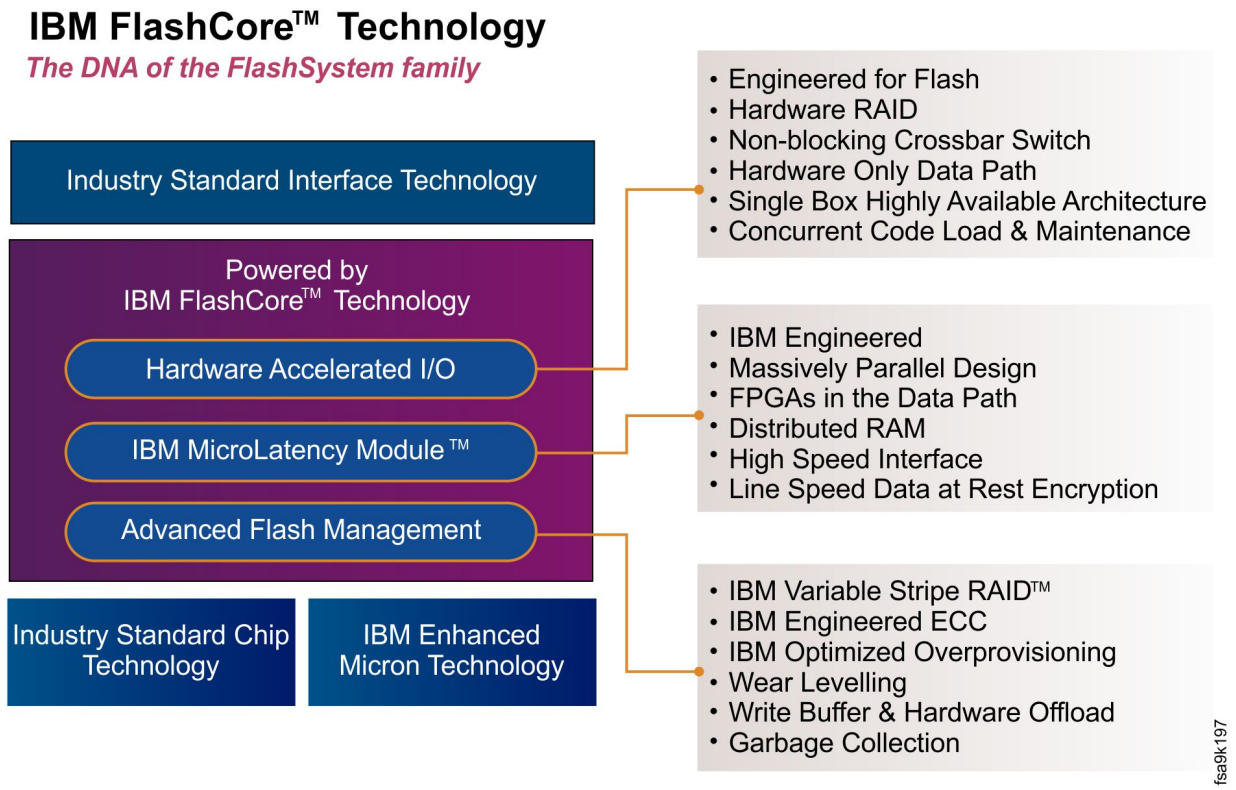

Grid controllers

Each grid controller is a high-end Intel Xeon-based compute server, which acts as a core component that provides the interface and advanced compute functions.

The grid controllers also provide cache to accelerate both read and write operations, and are responsible for carrying out the entire data reduction operation (see Flash-optimized data reduction). Data compression is provided by dedicated hardware accelerator cards.

Each grid controller is designed for high reliability and for modularity, so that components can be replaced in case of a failure, without interrupting the operation of the controller. By design, the grid controller is an isolated failure domain. Any failure or a maintenance action that requires shutdown of a grid controller, does not affect either the overall system operation or the protection status.

IBM FlashSystem® A9000R storage systems with enhanced grid controllers feature FC-NVMe ready adapters (see the Fibre Channel NVMe connectivity hardware announcement). NVM Express® (NVMe) allows servers to leverage the native parallelism of today's SSD offerings, reduces overall I/O overhead, and increases bandwidth. FC-NVMe enables NVMe over a Fibre Channel (FC) network fabric, thus combining the benefits of all-flash SAN storage with NVMe performance over existing infrastructure. A system is FC-NVMe ready if it requires a future software update to provide full FC-NVMe support.

To summarize, each grid controller contains:

- 56 Gbps InfiniBand HCAs
- Host interface ports
- Two hot-swappable hard disk drives (HDDs) that contain the system microcode and store various system logs and events
- Two hot-swappable solid-state drives (SSDs) that are used as vault devices
- Two hot-swappable battery backup units (BBUs)

- Two hot-swappable power supplies
- Field-replaceable fans
- Field-replaceable DIMMs

Back-end interconnect

Internal communication (interconnect) between grid controllers or flash enclosures is carried over a 56 Gbps Fourteen Data Rate (FDR) InfiniBand network infrastructure, with full redundancy.

In FlashSystem A9000R, the grid controllers and the flash enclosures are connected through redundant InfiniBand switches, and redundant cabling. The InfiniBand switches are also linked to each other.

The following figure describes the redundant InfiniBand connectivity with four grid controllers and two flash enclosures as an example, with redundant InfiniBand switches and cabling.

Figure 1. InfiniBand switch connectivity in FlashSystem A9000R

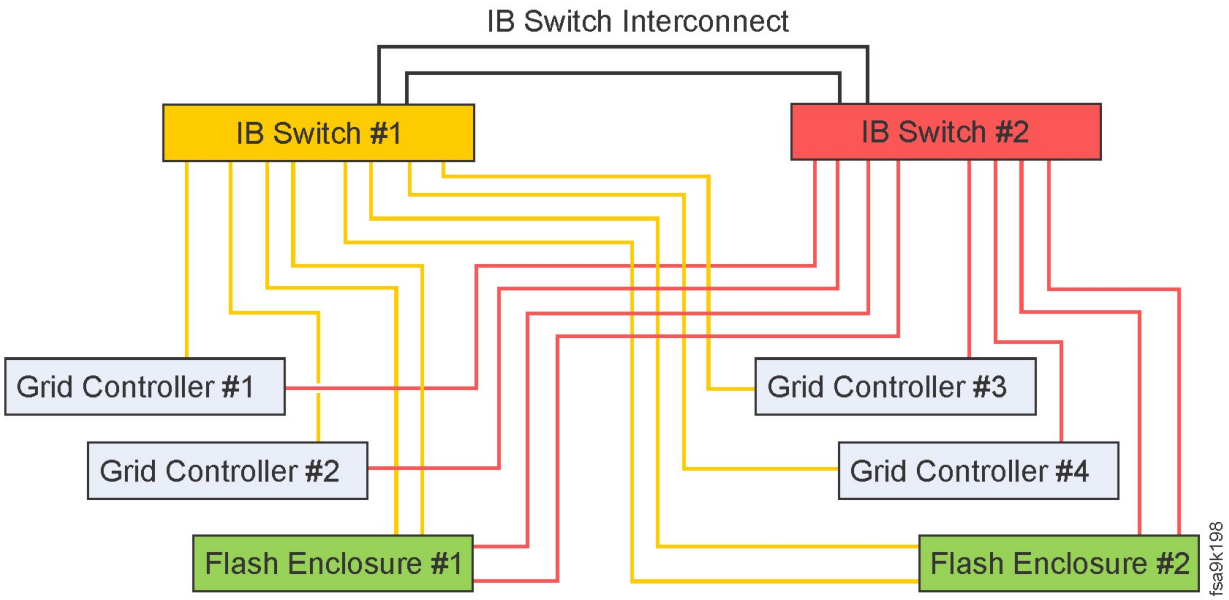

Logical architecture

The IBM FlashSystem® A9000R logical architecture is built on the IBM Spectrum Accelerate software, with many added enhancements to optimize the software stack for use with flash storage.

An important feature and differentiation in this new architecture is the software separation between the computation and storage resources of the product. In terms of the physical architecture, the compute functions are implemented in grid controllers, while the storage function is implemented in flash enclosures.

The separation of the cache, compute, and interface functions from the storage resource separates load balancing across compute and storage resources. It also enables support for a different resiliency scheme in which cache data is triplicated. This cache data protection is unique to IBM FlashSystem A9000R.

Another significant enhancement to the underlying Spectrum Accelerate software is the data reduction feature, that combines pattern matching and removal, deduplication, and compression. IBM FlashSystem A9000R also offers significant CPU processing power and RAM memory, which allows efficient processing of real-time (inline) deduplication and compression.

Scale-out grid architecture

The IBM FlashSystem® A9000R scale-out grid architecture enables the increase of the system's capacity and performance.

System scale-out is a non-disruptive process and can be done with any rack model that is not already fully populated. A scale-out is performed by adding additional grid controllers and flash enclosures to your system. For more information, see *IBM FlashSystem A9000R Deployment Guide* on the IBM FlashSystem A9000R documentation .

Functionality

IBM FlashSystem® A9000R delivers a range of advanced storage functional features.

These functional features include:

- Data reduction: pattern removal, deduplication, and compression
- Space-efficient snapshots and consistency groups
- Redirect-on-Write (ROW)
- Host rate limiting: Quality of Service (QoS) performance classes
- Synchronous remote mirroring
- Cross-generation asynchronous replication between A9000R systems and XIV® Gen3 systems
- HyperSwap®
- Multi-site HA/DR (high availability and disaster recovery)
- Volume migration with IBM® Hyper-Scale Mobility
- Multi-tenancy
- Data-at-rest encryption
- Remote support and proactive support
- Advanced statistics, including:
  - Capacity statistics
  - System performance metrics, including history metrics of no less than 300 days

The following figure illustrates how the first three functional features (see list above) are working in concert to provide top-of-the-line capabilities in an all-flash storage system.

Figure 1. Data reduction and protection

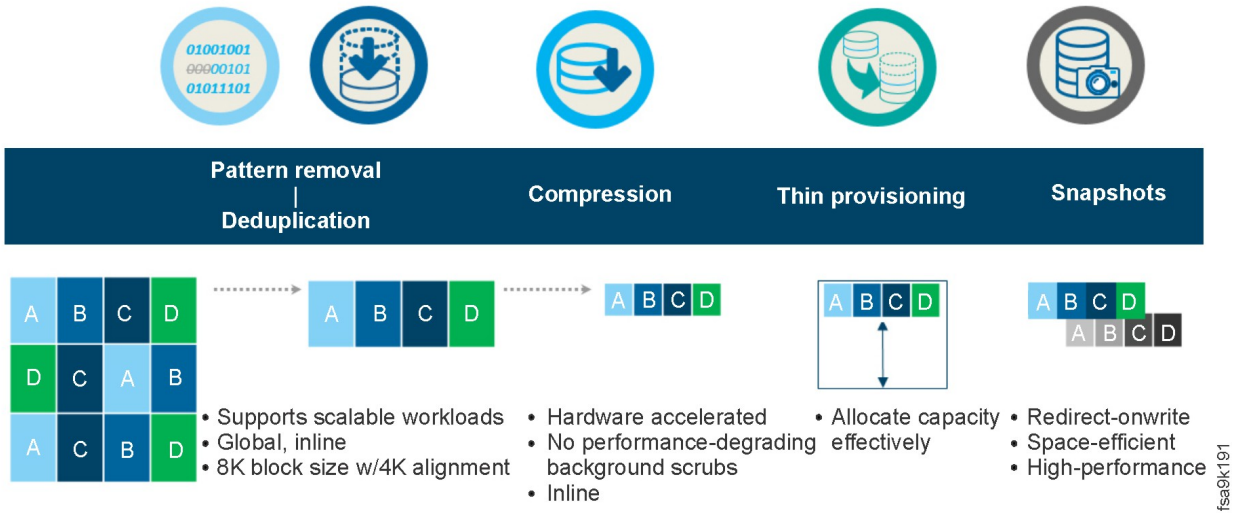

Specifications

IBM FlashSystem® A9000R hardware and other specifications, refer to the marketing data sheet or to the IBM FlashSystem A9000R Deployment Guide.

For specifications listed in the Deployment Guide, see fs_rack_dg_system_specs.html.
Note: For information about specification-related functional boundaries, such as maximum number of volumes and hosts per storage system, refer to the IBM FlashSystem A9000R specification sheet document, which can be provided to IBM® customers under non-disclosure agreement (NDA).

Performance

IBM FlashSystem® A9000R provides dynamic performance and ultra-low latency of IBM FlashCore® technology, integrated with an innovative grid architecture, extensive scalability, and industry-leading IBM® software-defined storage capabilities.

As a high-end all-flash storage platform, IBM FlashSystem A9000R is designed to address the most demanding requirements of enterprise clouds. The system architecture and range of grid components enable the system as a whole to deliver consistent microsecond latency with market-leading performance.

IBM Hyper-Scale Manager enables the monitoring of the accumulative duration of the internal and external system processing and data transport of an I/O request. For details, refer to IBM Hyper-Scale Manager documentation .

Flash-optimized data reduction

Ultra-fast and always-on data reduction is a fundamental key feature of the storage system.

The storage system uses industry-leading data reduction technology that combines in-line, real-time pattern matching and removal, deduplication, and compression. Together and in concert, these data reduction mechanisms allow the total physical storage capacity to be a few times larger, without affecting

performance even in heavy workloads.

The average data reduction ratio is approximately 5:1. However, the actual ratio could vary depending on the data type, as detailed in the following table.

Table 1. Data reduction ratio for different workloads

Workload type	Data reduction ratio
Virtual desktop infrastructure (VDI)	48:1 (98% savings)
VMware, Linux, and Windows	4:1 (74% savings)
Kernel-based virtual machine (KVM) on Linux	9:1 (89% savings)
Kernel-based virtual machine (KVM) on Windows	2:1 (55% savings)
Database	4:1 (77% savings)

Note: Other types of data environments and workloads may yield different reduction ratios.
Data reduction is implemented below the global cache to ensure very rapid response times, provide a global scope for data reduction services, and allow other data services to be completely unaffected, including snapshots, replication, offload-to-host features, and more.

With constant and automatic data reduction, the storage system's allocation limit is always larger than its raw physical capacity.

The following subsections describe the data reduction process and individual stages.

- **Data reduction stages**
  Each grid controller in the system dedicates some CPU processing capacity and memory resources for the purpose of data reduction.
- **Pattern detection and removal**
  Pattern detection and removal is the first stage in the data reduction process.
- **Inline deduplication**
  The inline deduplication is the second data reduction stage that processes and consolidates data before it is written to disk.
- **Inline compression**
  Inline compression is the last data reduction stage.

Data reduction stages

Each grid controller in the system dedicates some CPU processing capacity and memory resources for the purpose of data reduction.

The on-the-fly data reduction stages are:

1. Pattern detection and removal
2. Inline deduplication
3. Inline compression

Figure 1 illustrates how original data is reduced after each stage, before it is written to the flash disk.

Figure 1. Data reduction stages

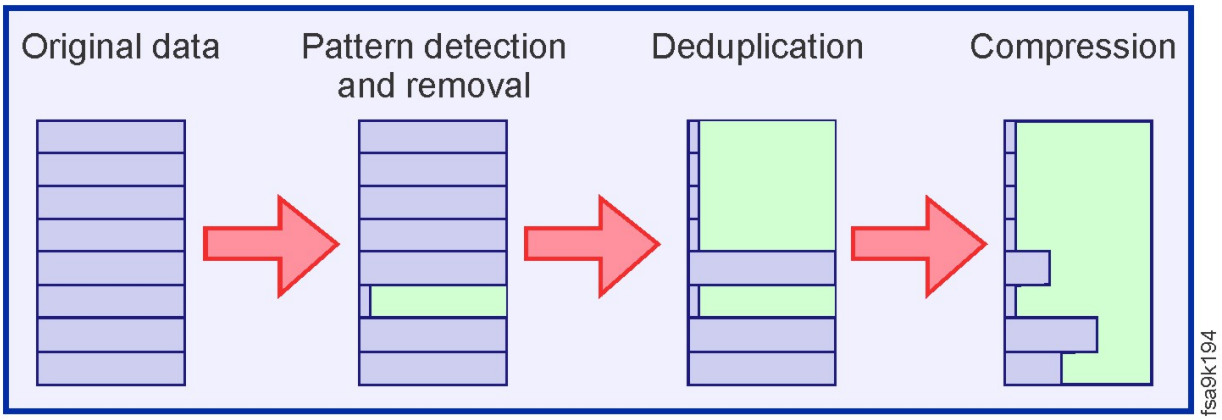

The system's data reduction node implements the data deduplication and compression functions.

The three data reduction stages are carried out according to a certain flow logic, as detailed in the flowchart below (Figure 2.

Figure 2. Data reduction process logic

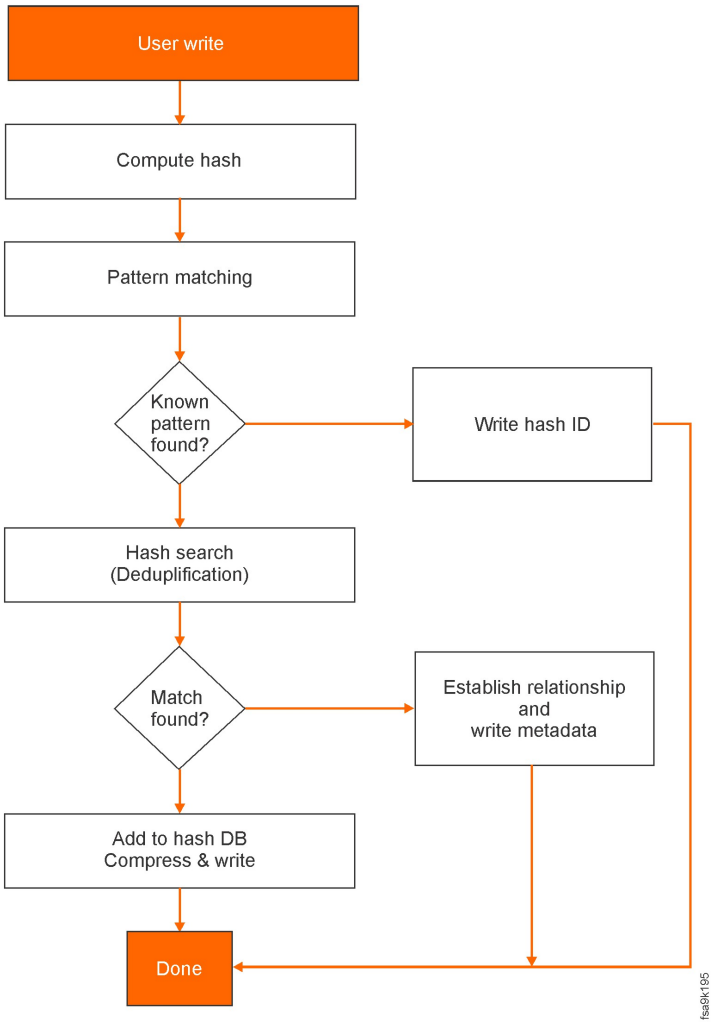

Pattern detection and removal

Pattern detection and removal is the first stage in the data reduction process.

Pattern matching mechanisms match incoming host writes with a preconfigured set of known patterns stored in the system. When a write is processed, it is split into 8 KB blocks, as shown in Figure 1.

Figure 1. Pattern detection and removal

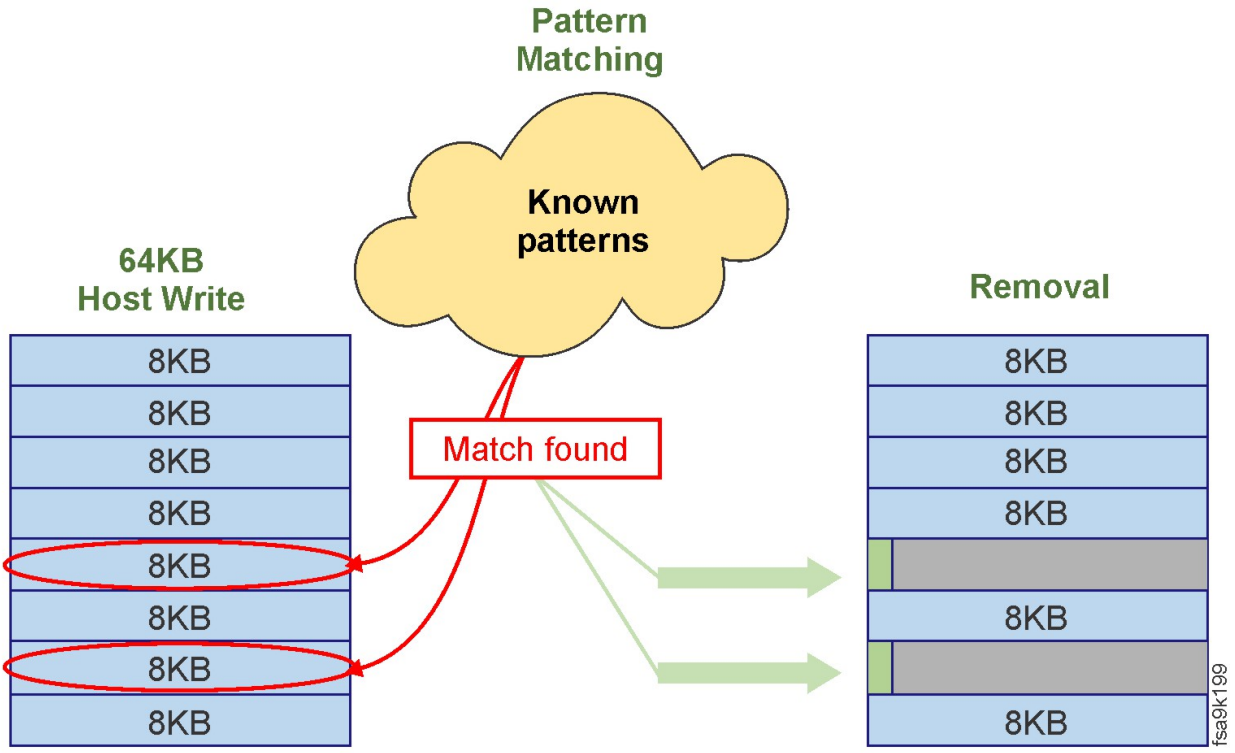

Each block is then hashed and the hash value, also known as a fingerprint, is compared to a table of well-known hashes. If a match is found, then the corresponding pattern ID, which is only 2 bytes is stored. Any match found at that stage is replaced with internal markings (a hash).

Inline deduplication

The inline deduplication is the second data reduction stage that processes and consolidates data before it is written to disk.

Hashing and hash comparison is performed on-the-fly. The benefit of inline deduplication is that duplicate chunks are never written to the destination disk system.

Figure 1 illustrates how original data is reduced via the deduplication process by using hashing and metadata.

Figure 1. Deduplication process

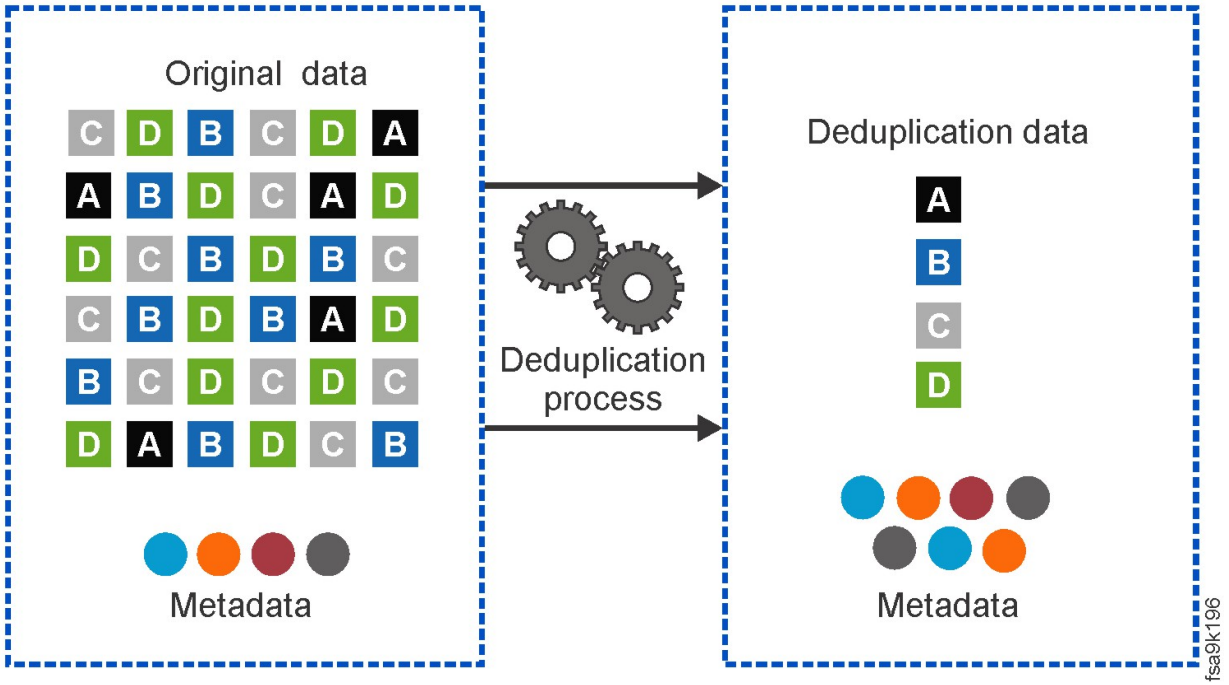

In this process:

- Hashes are calculated only once (upon write)
- Hashes are stored in metadata
- Data is compared by comparing the hash
- The chunk size is 8KB, although 4KB alignment is also supported and allows for a higher deduplication ratio

Data reduction with deduplication alone could reach to up to 30:1 (96.7% savings).

Inline compression

Inline compression is the last data reduction stage.

The IBM® patented compression technology used by the storage system is based on a data compression algorithm that operates in real time.

The major difference between traditional compression and IBM compression technique is in the size of data blocks that are written to the storage device. The IBM compression technique uses fixed-size writes (see Figure 1. This method enables an efficient and consistent method to index the compressed data, because it is stored in fixed-size containers.

Figure 1. IBM compression method - using fixed-size writes

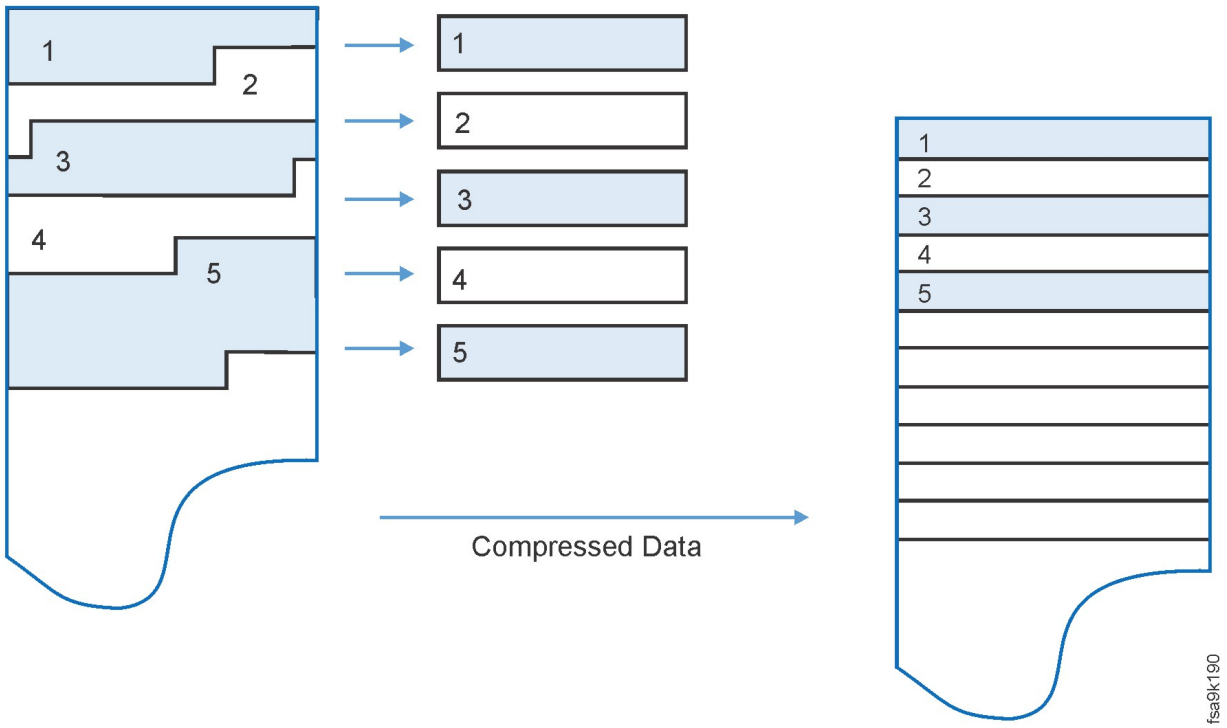

The compression ratio of a data block depends on how many repetitions can be detected within the block. The number of repetitions is affected by how many bytes present in the block are related to each other. The relation between bytes is driven by the format of the compressed object.

Flash-optimized data protection

The constant operation and availability of the system relies on the ratio of its functionality time to the total time it is required to function.

The system's high availability (HA) features pertain to both the hardware and software components. Hardware components are hot-swappable and support instant failover capabilities, allowing the system to have no single point of failure (SPOF).

In addition to 2D Flash RAID protection and Variable Stripe RAID data protection (see Two-dimensional flash RAID), the system's flash enclosure incorporates additional reliability features:

- Error correction codes to provide bit-level reconstruction of data from flash chips (see Scrubbing mechanism).
- Checksum and data integrity fields designed to protect all internal data transfers within the system.
- Overprovisioning to enhance write endurance and decrease write amplification.
- Wear-leveling algorithms balance the number of writes among flash chips throughout the system.
- Sweeper algorithms help ensure that all data within the system is read periodically to avoid data fade issues.

Additional system protection mechanisms include:

- Advanced vaulting mechanism, including live vaulting (see Vaulting mechanism).
- All cached data is protected using triplication (3 copies).
- Software nodes are redundant with auto-restart of hanging nodes.
- Any unresponsive data node is expelled from the cluster automatically.

- Upon any performance degradation, events are issued and reported.
- System temperature monitoring with warning events generation.
- Power monitoring, including battery unit conditioning (calibration) for increased capacity and lifespan.

- **Two-dimensional flash RAID**
  The combination of IBM Variable Stripe RAID and system-level RAID 5 protection across IBM MicroLatency modules (within a flash enclosure) is called two-dimensional (2D) flash RAID.
- **Vaulting mechanism**
  Each grid controller contains two 400 GB enterprise-grade SSD disks for saving cache data in case of power loss, in a process called *vaulting*.
- **Scrubbing mechanism**
  Scrubbing is a background process that systematically verifies there are no data inconsistencies, and, if possible, amends the detected errors.

Two-dimensional flash RAID

The combination of IBM® Variable Stripe RAID and system-level RAID 5 protection across IBM MicroLatency® modules (within a flash enclosure) is called two-dimensional (2D) flash RAID.

Two-dimensional (2D) flash RAID consists of IBM Variable Stripe RAID and system-wide RAID 5.

Variable Stripe RAID technology (see Figure 1) helps reduce downtime and maintain performance and capacity in the event of partial or full flash chip failures. Failures of large capacity granularity, such as an entire flash chip or a plane or a die of a flash chip, are isolated, reduced to the absolute minimum affected area, then bypassed.

Figure 1. Variable Stripe RAID technology

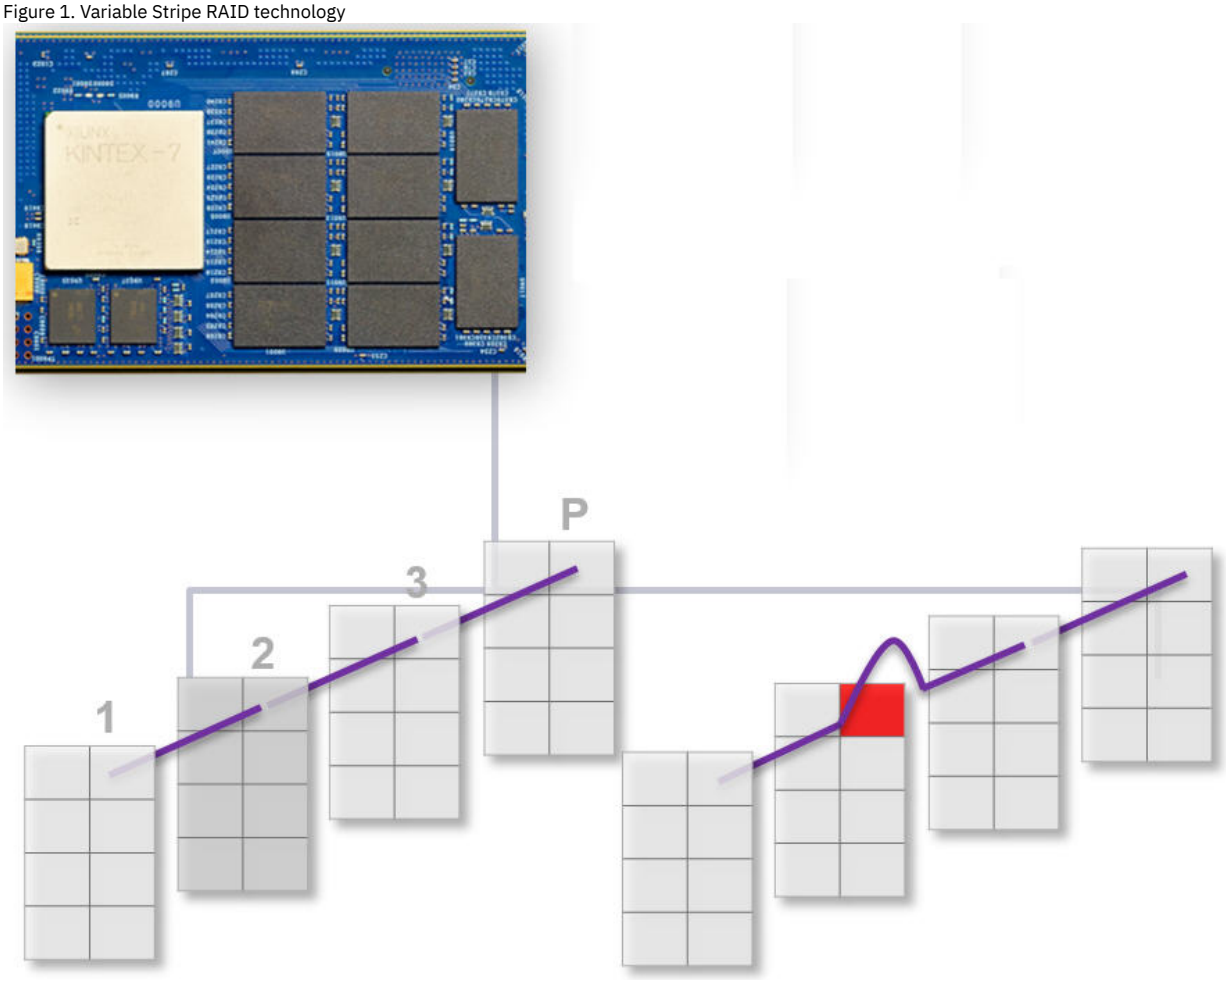

System-wide RAID 5, with easily accessed hot swappable flash modules, helps promote availability. RAID 5 configurations (Figure 2) provide a high degree of redundancy with Variable Stripe RAID and RAID 5 protection. RAID 5 data protection includes one IBM MicroLatency module dedicated as parity and another as a dedicated hot spare. The maximum capacity utilization for RAID 5 is provided by using 12 IBM MicroLatency modules.

Figure 2. 2D RAID protection mechanism

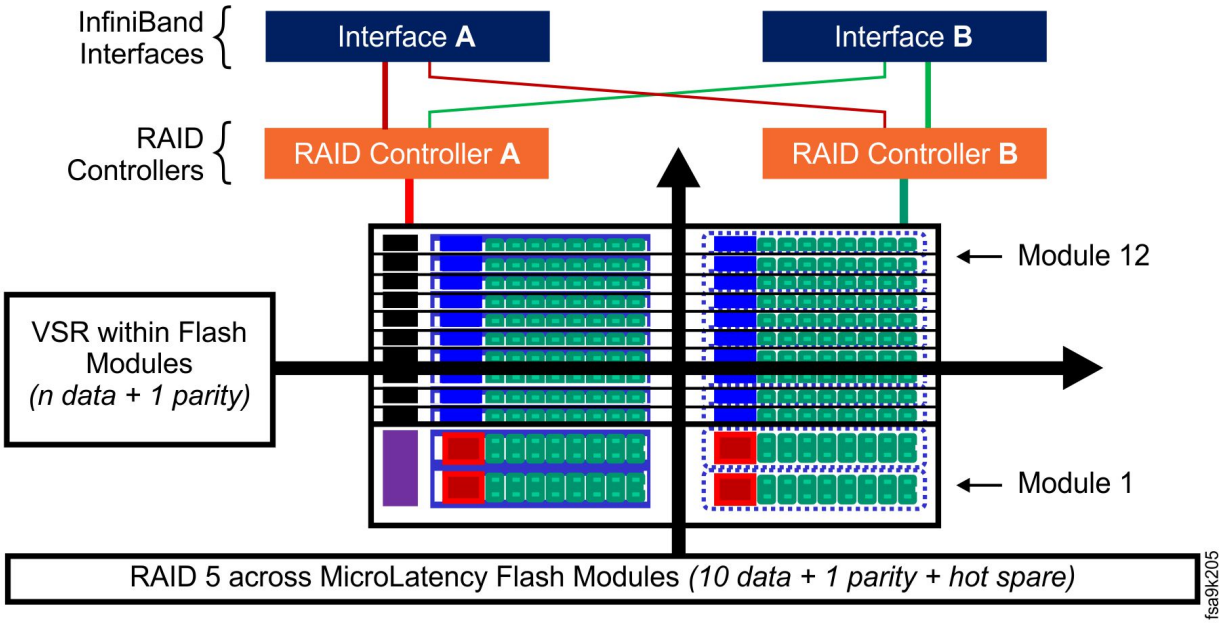

The system's flash enclosure level RAID 5 complements the Variable Stripe RAID technology implemented within each MicroLatency module, providing protection against data unavailability resulting from MicroLatency module failures. It also allows data to be rebuilt onto a hot spare flash module, so that any individual MicroLatency module could be replaced without data disruption.

Vaulting mechanism

Each grid controller contains two 400 GB enterprise-grade SSD disks for saving cache data in case of power loss, in a process called *vaulting*.

In addition, the SSD disks are used to save metadata and system configuration information on a continuous basis, in a process called *live vaulting*. During a normal shutdown procedure, the system's microcode also writes to the SSD disks all configuration data, metadata, and cache data that has not yet been stored in the permanent flash enclosure storage.

Figure 1 shows how the cache data is mirrored three times across different grid controllers, without requiring mirroring within the grid controller.

Figure 1. 3 copies of cache data

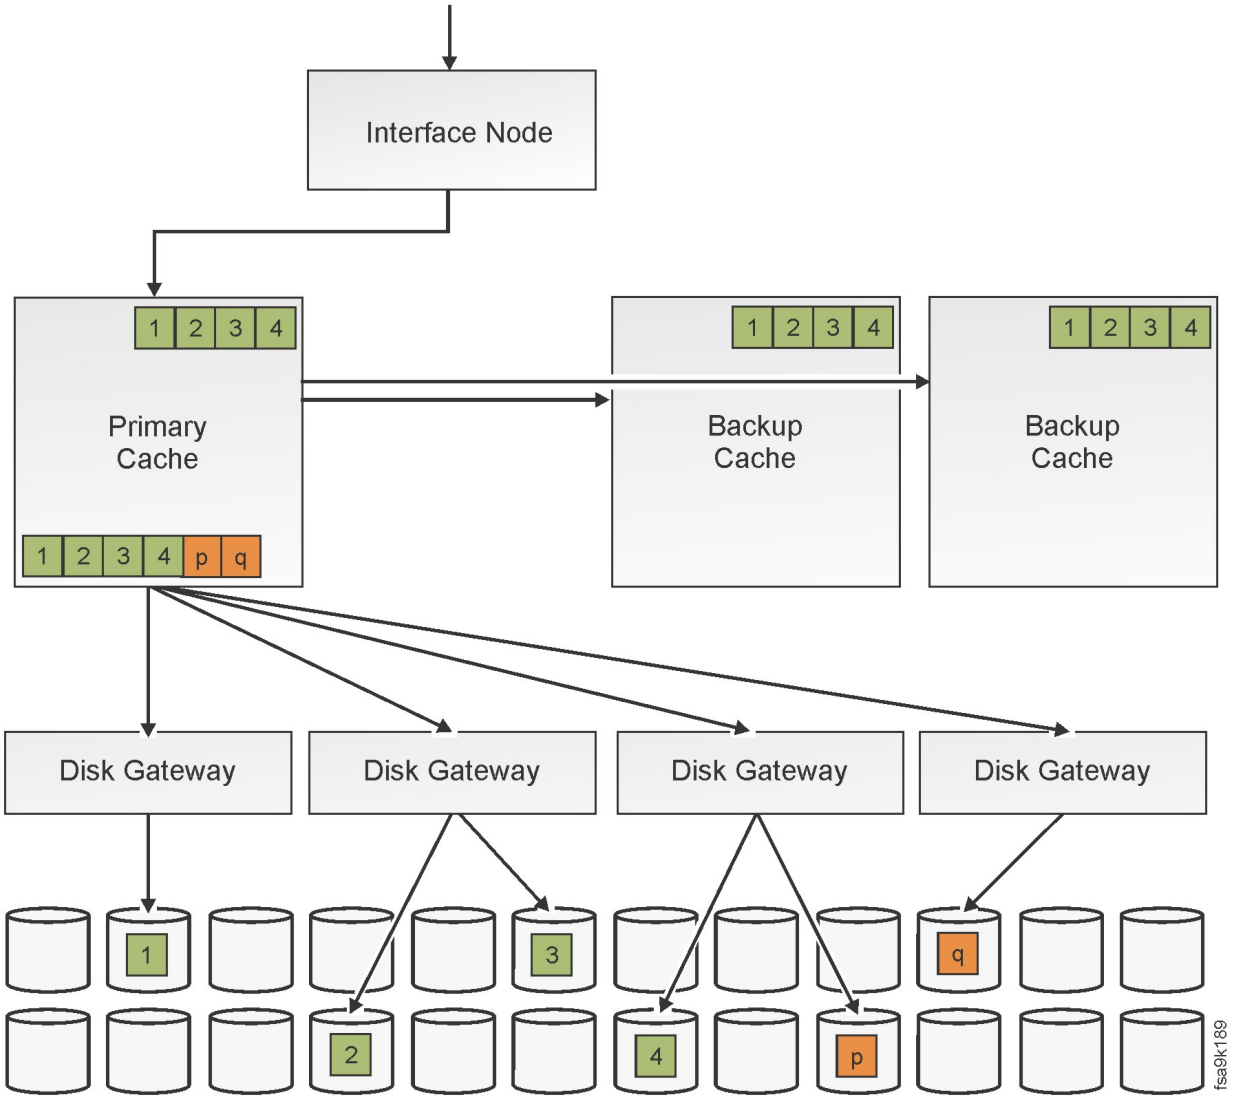

Scrubbing mechanism

Scrubbing is a background process that systematically verifies there are no data inconsistencies, and, if possible, amends the detected errors.

Physical data scrubbing is performed autonomously and internally by the flash enclosure. The scrubbing process verifies the consistency of the RAID 5. Error correction is performed for recoverable errors (single error), and no correction is performed for unrecoverable errors (double error).

For unrecoverable errors, an event is generated and reported, but the flash enclosure remains up and running.

Physical scrubbing is complemented by virtual scrubbing which verifies the data through RACE. This additional scrubbing process can also make an effort to correct the detected errors (for example, by using mirroring). Virtual scrubbing, in its turn, is combined with remote scrubbing, which also sequentially reads the data through RACE. When the volume is primary in remote mirroring (see Synchronous remote mirroring), remote scrubbing compares the content on the partition to its counterpart on the remote target. If the partition contents are not identical, an event is issued.

Flash-optimized data path

The hardware-only data path design of the flash enclosure eliminates latency at the system's software layer. The system's data path is completely independent of the control path.

Data traverses through the flash enclosure through field-programmable gate arrays (FPGAs), preventing wasted cycles on interface translation, protocol control, or tiering. FPGAs are included in the interface cards, RAID controllers, and MicroLatency modules.

In addition, all data passes through three nodes:

- Interface nodes
- Cache nodes
- Data reduction nodes

These nodes reside on each grid controller. The system load is uniformly distributed between all the grid controllers, resulting with a statistically balanced load. All data nodes communicate with each other, so that data entering the system through a certain interface node can be directed to any of the cache nodes, in any of the grid controllers.

The data flow includes the following stages:

1. I/O requests enter the system as SCSI commands on any of the interface nodes running on the grid controllers. The I/O requests are forwarded to the cache nodes based on a distribution table, which maintains rules to ensure that the I/O and capacity loads are equally distributed.
2. The cache data is copied three times across different grid controllers, providing protection in case of any unexpected failure.
3. The cache node directs the data to any of the data reduction nodes on any of the grid controllers.
4. The data reduction node performs data reduction and then moves it back to one of the cache nodes.
5. The cache node then acts as a gateway, performing a read/write operation to the flash enclosure.

Figure 1 illustrates these stages.

Figure 1. Data path through the different nodes

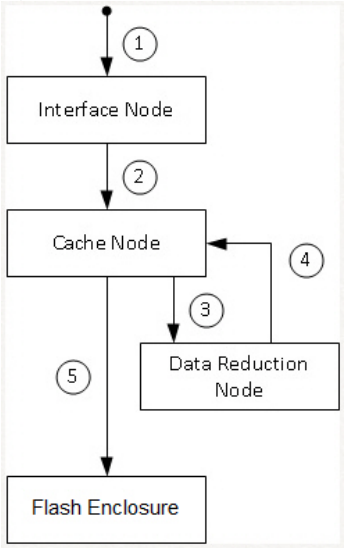

Upon any failure in one of the grid controllers, or upon detection of a slow node, the system automatically and seamlessly re-adjusts the load distribution, and the data of the failed/slow component is shared among the rest of the nodes, to avoid performance degradation and to re-balance the system.

FlashSystem A9000 systems include only one flash enclosure. However, on FlashSystem A9000R systems that include more than one flash enclosure, data can be moved from any grid controller to any flash enclosure. This mesh architecture enables the system to evenly distribute workloads in the system, as shown in Figure 2.

Figure 2. Data path mesh architecture

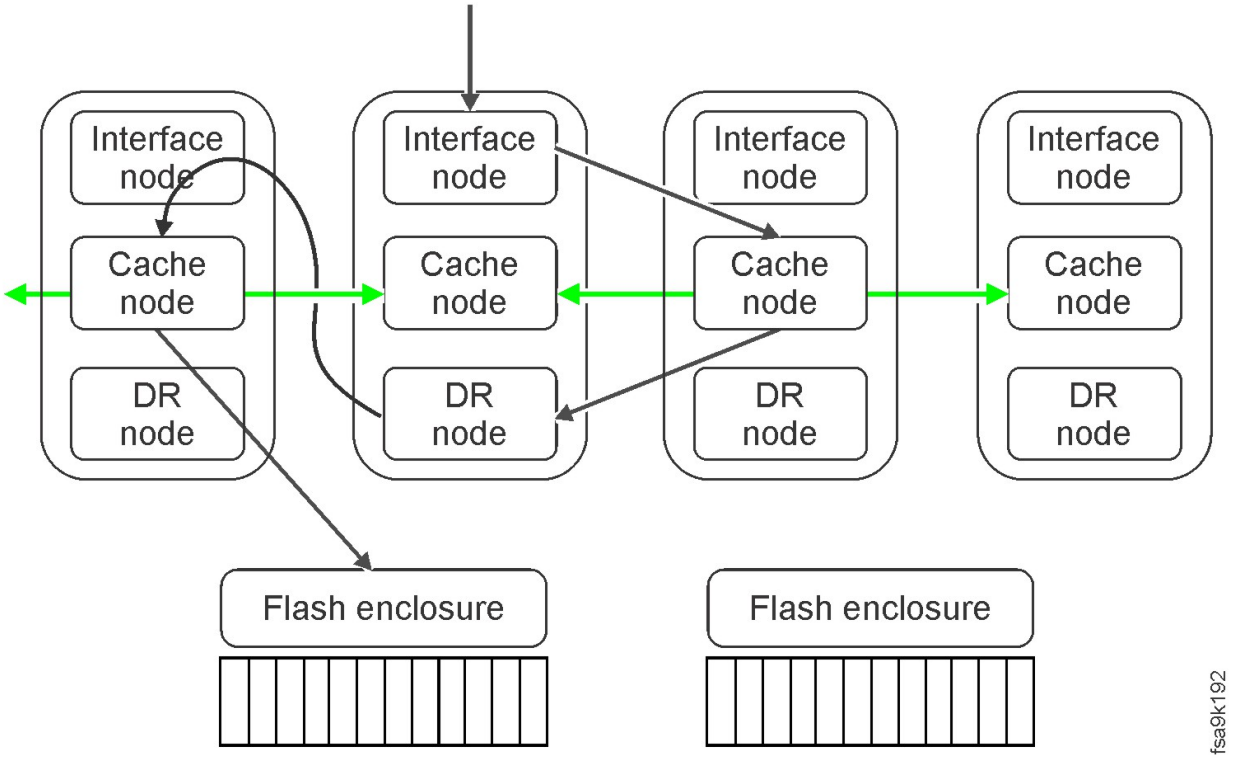

Capacity management

Managing and utilizing the storage system's capacity is a key aspect of the system management.

Capacity management hierarchy in a storage system is shown in the diagram below:

Figure 1. Capacity management hierarchy

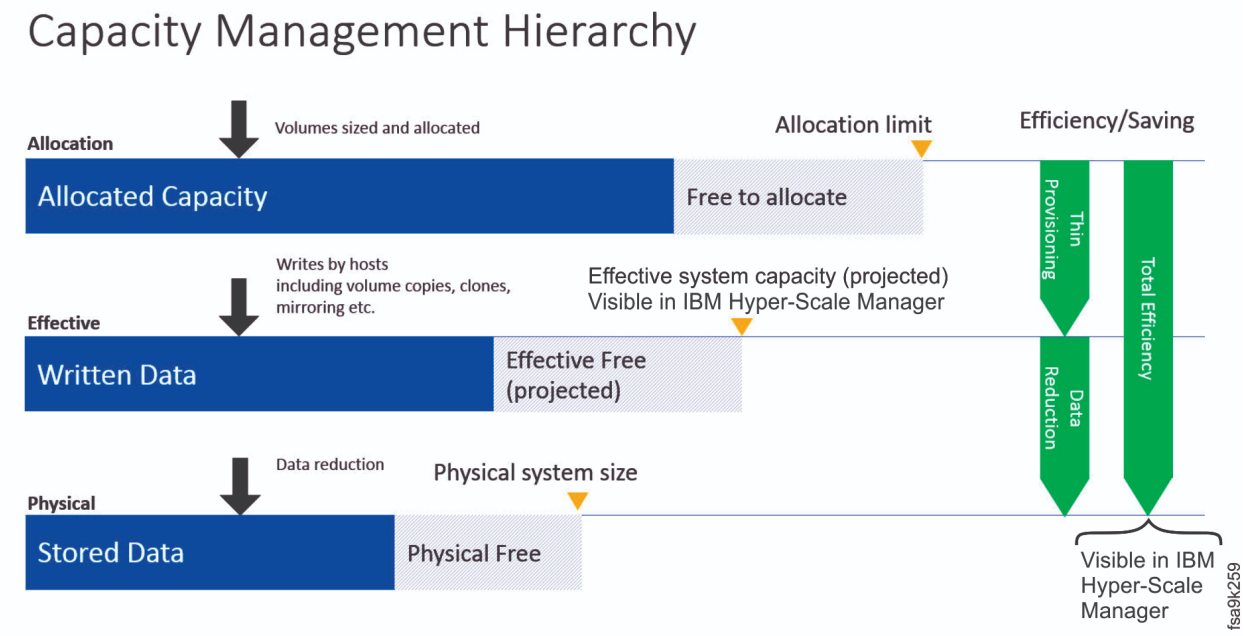

- **Physical size** – Represents the physical amount of storage space that is available in the system. In systems with more than a single flash enclosure, the physical capacity is the sum of capacity of all flash enclosures combined. When a volume or a snapshot is deleted, the physical size value in the system capacity monitoring tools (IBM Hyper-Scale Manager and CLI commands) is updated not immediately, but gradually. The amount of physical storage to be reclaimed appears in the system capacity monitoring tools as *Pending Deletion*.
- **Effective capacity** – Represents the total data written by hosts to volumes and snapshots, before applying data reduction (see Flash-optimized data reduction). This includes volume copies, clones, mirroring, and so on.
- **Allocation limit** – Represents the actual available-for-use capacity after data reduction. The allocation limit varies depending on the type of data being written, and the data reduction ratio that is achieved for that type of data. With constant and automatic data reduction, the allocation limit is always much larger than the system's raw physical capacity. The segment of the allocation limit that is currently allocated to volumes and snapshots is referred to as **allocated capacity** (see also Allocation limit).

Capacity management includes the following aspects:

- Provisioning – Defining resources and assigning these resources to storage objects (domains, storage pools, volumes, snapshots).
- Accounting – Tracking resource consumption at various levels, and making sure that provisioned boundaries are observed.
- Statistics – Collecting capacity usage statistics over different periods of time.
- Monitoring – Defining capacity usage thresholds and alerts.

When managing the system capacity fromIBM® Hyper-Scale Manager, additional projection and planning capabilities and management options are available. For details, see IBM Hyper-Scale Manager documentation .
Note: Customers looking for pay-per-use capacity licensing of FlashSystem A9000 or A9000R are advised to consider FlashSystem A9000 or A9000R model U25, a Storage Utility Offering that is based on the same hardware as model 425. For details, refer to the IBM Storage Utility Offering home page (ibm.com/products/storage-utility).

- **Thin provisioning**
  Thin provisioning, or over-provisioning, provides the ability to define logical volume or storage pool sizes that are much larger than the physical capacity of the system.
- **Allocation limit**
  Increase in the physical capacity of a single storage system or a group of scaled-out storage systems results in an increase of the maximum allocation limit.
- **Deduplication**
  Deduplication increases storage efficiency by removing duplicate data.

Thin provisioning

Thin provisioning, or over-provisioning, provides the ability to define logical volume or storage pool sizes that are much larger than the physical capacity of the system.

Capacity associated with specific applications or users can be dynamically increased or decreased per demand, without necessitating an accurate prediction of future needs. Physical capacity is only committed when the associated applications execute writes, and not when the logical volume is initially allocated. In other words, thin provisioning is the virtualization of the capacity from the underlying hardware allocation.

Because the total system capacity is designed as a globally available pool, thin provisioned resources share the “buffer” of free space. This approach results in highly efficient aggregate capacity use without pockets of inaccessible unused space.

However, because the storage system already implements data reduction that is carried out automatically (see Flash-optimized data reduction), thin provisioning is provided by default, because the allocation limit is already at least 5 times greater than the physical capacity of the system.

The allocatable capacity in the system can be administratively portioned into separate and independent storage pools (see Storage pools). The storage administrator can create volumes (along with the reserved snapshot space) whose total capacity is less than or equal to the size of the pool. At the same time, the data reduction methods of deduplication and compression are reducing the amount of data actually written to the flash enclosure.

Allocation limit

Increase in the physical capacity of a single storage system or a group of scaled-out storage systems results in an increase of the maximum allocation limit.

Physical capacity addition is followed by a process of redistribution (restriping) of the existing data across the bigger physical capacity. Only at the end of this phase-in process will the physical capacity be updated to reflect the expanded physical and effective capacity.

The allocation limits for different configurations of IBM FlashSystem® A9000R are listed in the table below.

Table 1. Allocation limits for various IBM FlashSystem A9000R configurations

Model	Number of grid controllers	Allocation limit	Number of grid elements
415	4	1400 TB	2
	6	2000 TB	3
	8	2600 TB	4
	10 and more	3000 TB	5 and more
425 and U25	3	1200	1 and 1 additional grid controller
	4	2400 TB	2
	6	3600 TB	3
	8	4800 TB	4

Deduplication

Deduplication increases storage efficiency by removing duplicate data.

Deduplication-aware capacity management reports actual and more granular estimates of physical and reclaimable capacity per volume, as well as estimates of deduplication and compression ratios. The deduplication-aware capacity management calculations are built to support future optimization tasks, such as determining the best system – from among multiple systems – to place a volume. For more information, see IBM® Hyper-Scale Manager documentation .

Volumes and snapshots

Volumes are the basic storage data units in the storage system. Snapshots of volumes can be created, where a snapshot of a volume represents the data on that volume at a specific point in time.

Volumes can also be grouped into larger sets called consistency groups (see Consistency groups) and storage pools (see Storage pools).

The basic hierarchy is as follows:

- A volume can have multiple snapshots.
- A volume can be part of one and only one consistency group.
- A volume is always a part of one and only one storage pool.
- All volumes in a consistency group must belong to the same storage pool.

- **Volume function and lifecycle**
  Volume is the basic data container that is presented to the hosts as a logical disk.
- **Snapshot function and lifecycle**
  A *snapshot* is a logical volume reflecting the contents of a given source volume at a specific point-in-time.
- **Additional snapshot attributes**
  Snapshots have the following additional attributes.
- **Redirect-on-Write (ROW)**
  The storage system uses the Redirect-on-Write (ROW) mechanism.
- **Full volume copy**
  Full Volume Copy overwrites an existing volume, and at the time of its creation it is logically equivalent to the source volume.
- **Restoring volumes and snapshots**
  The restoration operation provides the user with the ability to instantly recover the data of a Primary volume from any of its locked snapshots.

Volume function and lifecycle

Volume is the basic data container that is presented to the hosts as a logical disk.

The term *volume* is sometimes used for an entity that is either a volume or a snapshot. Hosts view volumes and snapshots through the same protocol. Whenever required, the term *Primary volume* is used for a volume to clearly distinguish volumes from snapshots.

Each volume has two configuration attributes: a name and a size. The volume name is an alphanumeric string that is internal to the storage system and is used to identify the volume to both the GUI and CLI commands. The volume name is not related to the SCSI protocol. The volume size represents the number of blocks in the volume that the using host detects.

A volume is defined within the context of only one storage pool (see Storage pools). Because storage pools are logical constructs, a volume and any snapshots associated with it can be moved to any other storage pool (within the same domain) if there is sufficient space within the target storage pool.

Volumes in FlashSystem A9000 and A9000R

As a benefit of the system virtualization, there are no limitations on the associations between logical volumes and storage pools. Moreover, manipulation of storage pools consists exclusively of metadata transactions and does not trigger any copying of data. Therefore, changes are completed instantly and without any system performance degradation.

The storage system uses the grid concept and distributes volume data evenly across hardware storage resources. Volumes are distributed evenly across all flash enclosures using partitions and each partition is 16 MB in size.

The system also uses the concept of allocation unit (AU) size for volumes, which is set at 103 GB. Minimum volume size that can be created is 1 GB. However, volumes that are created with a specified size about five percent or less, smaller than the AU size or multiples of allocation unit size will be rounded to multiples of AU size. For example, creating a volume, specifying a 98 GB size creates a volume of 103 GB on the system.

Volume management and lifecycle

The following management options are available for volumes. Volume operation is shown in Figure 1.

Create
: Defines the volume using the attributes you specify.

Resize
: Changes the virtual capacity of the volume.

Copy
: Copies the volume to an existing volume or to a new volume.

Format
: Clears the volume.

Lock
: Prevents hosts from writing to the volume.

Unlock
: Allows hosts to write to the volume.

Rename
: Changes the name of the volume, while maintaining all of the volumes previously defined attributes.

Delete
: Deletes the volume.

Figure 1. Volume operations

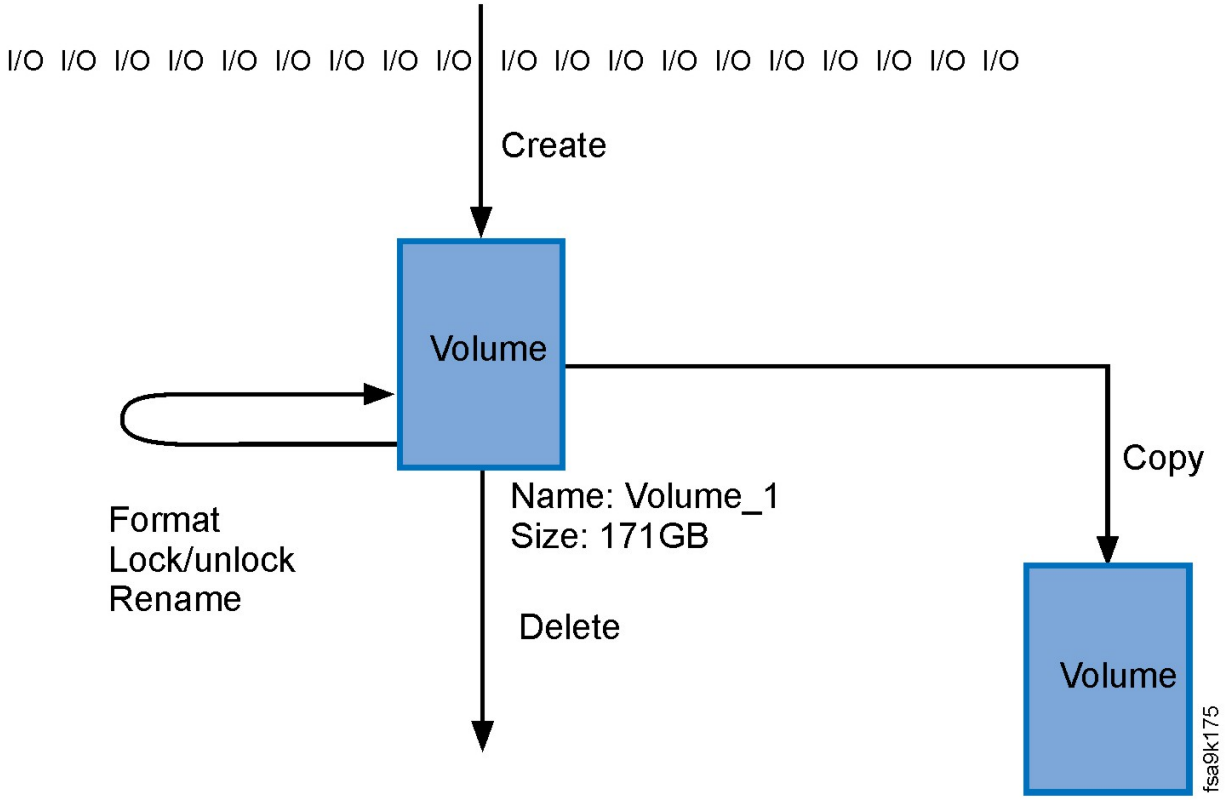

The following query options are available:

Listing volumes
: List the details of all volumes, or a specific volume according to a given volume or pool.

Finding a volume based on a SCSI serial number
: Display the volume name according to its SCSI serial number.

Snapshot function and lifecycle

A *snapshot* is a logical volume reflecting the contents of a given source volume at a specific point-in-time.

The storage system uses advanced snapshot mechanisms to create a virtually unlimited number of volume copies without impacting performance. Snapshot taking and management are based on a mechanism of internal pointers that allow the Primary volume and its snapshots to use a single copy of data for all portions that have not been modified.

This approach, also known as Redirect-on-Write (ROW) is an improvement of the more common Copy-on-Write (COW), which translates into a reduction of I/O actions, and therefore storage usage. For more information, see Redirect-on-Write (ROW).

No storage capacity is consumed by the snapshot until the source volume (or the snapshot) is changed.

Figure 1 shows the snapshot's lifecycle and the various operations available to the user.

Figure 1. The snapshot life cycle

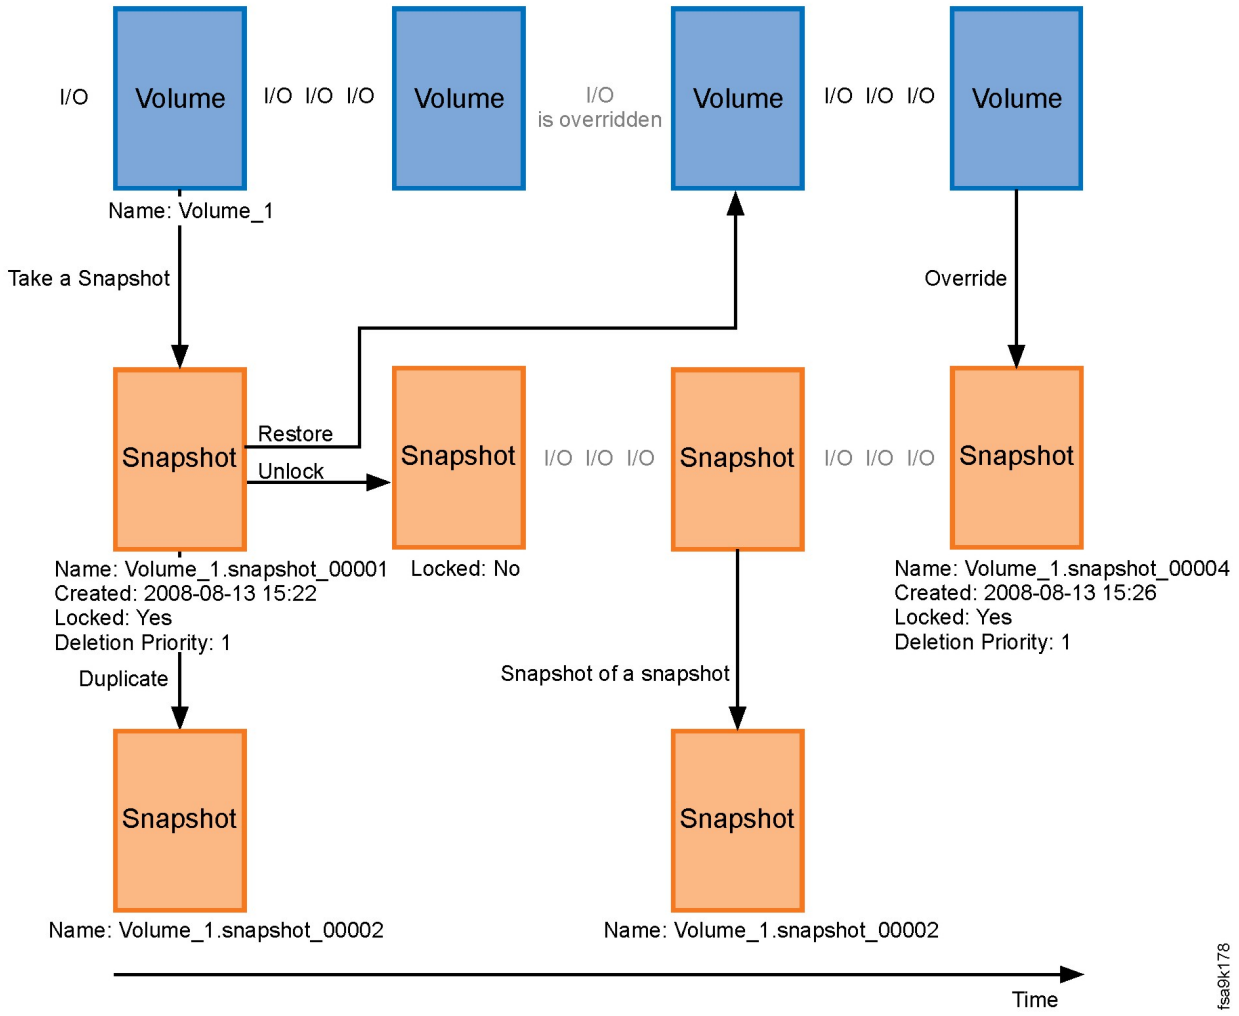

The following operations are applicable for the snapshot:

Create
: Creates the snapshot (taking a snapshot).

Restore
: Copies the snapshot back onto the volume. The main snapshot functionality is the capability to restore the volume.

Unlocking
: Unlocks the snapshot to make it writable and sets the status to Modified. Re-locking the unlocked snapshot disables further writing, but does not change the status from Modified.

Duplicate
: Duplicates the snapshot. Similar to the volume, which can be snapshotted infinitely, the snapshot itself can be duplicated.

A snapshot of a snapshot
: Creates a backup of a snapshot that was written into. Taking a snapshot of a writable snapshot is similar to taking a snapshot of a volume.

Overwriting a snapshot
: Overwrites a specific snapshot with the content of the volume.

Delete
: Deletes the snapshot.

The following subsections provide additional details about the snapshot lifecycle stages.

- **Creating a snapshot**
  First, a snapshot of the volume is taken.
- **Locking and unlocking a snapshot**
  Initially, a snapshot is created in a locked state, which prevents it from being changed in any way related to data or size, and only enables the reading of its contents.
- **Duplicating a snapshot**
  Authorized users can create a new snapshot by duplicating an existing snapshot.
- **Creating a snapshot of a snapshot**
  When duplicating a snapshot that was changed using the unlock feature, the generated snapshot is actually a snapshot of a snapshot.
- **Formatting a snapshot or a snapshot group**
  The format operation deletes the content of a snapshot - or a snapshot group - while maintaining its mapping to the host.

Creating a snapshot

First, a snapshot of the volume is taken.

The system creates a pointer to the volume, hence the snapshot is considered to have been immediately created. This is an atomic procedure that is completed in a negligible amount of time. At this point, all data portions that are associated with the volume are also associated with the snapshot.

Later, when a request arrives to read a certain data portion from either the volume or the snapshot, it reads from the same single, physical copy of that data.

When a write I/O request arrives to the Primary volume, a data reference to the overwritten data is created and associated with the Primary volume snapshot. This way, the overwritten data is retained in the snapshot.
Note: For blocks smaller than 8 KB, a copy of the overwritten data is created and associated with the Primary volume snapshot instead.

Locking and unlocking a snapshot

Initially, a snapshot is created in a locked state, which prevents it from being changed in any way related to data or size, and only enables the reading of its contents.

This is called an *image* or *image snapshot* and represents an exact replica of the Primary volume when the snapshot was created.

A snapshot can be unlocked after it is created. The first time a snapshot is unlocked, the system initiates an irreversible procedure that puts the snapshot in a state where it acts like a regular volume with respect to all changing operations. Specifically, it allows write requests to the snapshot. This state is immediately set by the system and brands the snapshot with a permanent modified status, even if no modifications were performed. A *modified snapshot* is no longer an image snapshot.

An unlocked snapshot is recognized by the hosts as any other writable volume. It is possible to change the content of unlocked snapshots, however, physical storage space is consumed only for the changes. It is also possible to resize an unlocked snapshot.

Primary volumes can also be locked and unlocked. A locked Primary volume cannot accept write commands from hosts. The size of locked volumes cannot be modified.

Duplicating a snapshot

Authorized users can create a new snapshot by duplicating an existing snapshot.

The snapshot duplicate is identical to the source snapshot. The new snapshot is associated with the Primary volume of the existing snapshot, and appears as if it were taken at the exact moment the source snapshot was taken. For image snapshots that have never been unlocked, the duplicate is given the exact same creation date as the original snapshot, rather than the duplication creation date.

With this feature, a user can create two or more identical copies of a snapshot for backup purposes, and perform modification operations on one of them without sacrificing the usage of the snapshot as an untouched backup of the Primary volume, or the ability to restore from the snapshot.

Creating a snapshot of a snapshot

When duplicating a snapshot that was changed using the unlock feature, the generated snapshot is actually a snapshot of a snapshot.

The creation time of the newly created snapshot is when the command was issued, and its content reflects the contents of the source snapshot at the moment of creation.

After its creation, the new snapshot is viewed as another snapshot of the Primary volume.

Formatting a snapshot or a snapshot group

The format operation deletes the content of a snapshot - or a snapshot group - while maintaining its mapping to the host.

The purpose of the formatting is to allow customers to back up their volumes via snapshots, while maintaining the snapshot ID and the LUN ID. More than a single snapshot can be formatted per volume.

Format operation results

The format operation results with the following:

- The formatted snapshot is read-only
- The format operation has no impact on performance
- The formatted snapshot does not consume space
- Reading from the formatted snapshot always returns zeroes
- It can be overridden
- It can be deleted
- Its deletion priority can be changed

Restrictions

No unlock
: The formatted snapshot is read-only and cannot be unlocked.

No volume restore
: The volume that the formatted snapshot belongs to cannot be restored from it.

No restore from another snapshot
: The formatted snapshot cannot be restored from another snapshot.

No duplicating
: The formatted snapshot cannot be duplicated.

No re-format
: The formatted snapshot cannot be formatted again.

No volume copy
: The formatted snapshot cannot serve as a basis for volume copy.

No resize
: The formatted snapshot cannot be resized.

Snapshot formatting example

1. Create a snapshot for each LUN you would like to back up to, and mount it to the host.
2. Configure the host to back up this LUN.
3. Format the snapshot.
4. Re-snap. The LUN ID, snapshot ID, and host mapping are maintained.

Restrictions in relation to other operations

Snapshots of the following types cannot be formatted:

Internal snapshot
: Formatting an internal snapshot hampers the process it is part of.

Part of a sync job
: Formatting a snapshot that is part of a sync job renders the sync job meaningless.

Part of a snapshot group
: A snapshot that is part of a snapshot group cannot be treated as an individual snapshot.

Snapshot group restrictions
: All snapshot format restrictions apply to the snapshot group format operation.

Additional snapshot attributes

Snapshots have the following additional attributes.

Storage utilization

The storage system allocates space for volumes and their snapshots in a way that whenever a snapshot is taken, additional space is actually needed only when the volume is written into.

As long as there is no actual writing into the volume, the snapshot does not need actual space. However, some applications write into the volume whenever a snapshot is taken. This writing into the volume mandates immediate space allocation for this new snapshot. Hence, these applications use space less efficiently than other applications.

Auto-delete priority

Snapshots are associated with an *auto-delete priority* to control the order in which snapshots are automatically deleted.

Taking volume snapshots gradually fills up storage space according to the amount of data that is modified in either the volume or its snapshots. To free up space when the maximum storage capacity is reached, the system can refer to the auto-delete priority to determine the order in which snapshots are deleted. If snapshots have the same priority, the snapshot that was created first is deleted first.

Name and association

A snapshot can either be taken of a source volume, or from a source snapshot.

The name of a snapshot is either automatically assigned by the system at creation time or given as a parameter of the CLI command that creates it. The snapshot's auto-generated name is derived from its volume's name and a serial number. The same applies when using the management GUI.

The following are examples of snapshot names:

```
MASTERVOL.snapshot_XXXXX
NewDB-server2.snapshot_00597
```

Table 1. Example parameters for source snapshots

Parameter	Description	Example
MASTERVOL	The name of the volume.	NewDB-server2

Parameter	Description	Example
XXXXX	A five-digit, zero filled snapshot number.	00597

Redirect-on-Write (ROW)

The storage system uses the Redirect-on-Write (ROW) mechanism.

The following items are characteristics of using ROW when a write request is directed to the Primary volume:

1. The data originally associated with the Primary volume remains in place.
2. The new data is written to a different location on the disk.
3. After the write request is completed and acknowledged, the original data is associated with the snapshot and the newly written data is associated with the Primary volume.

In contrast with the traditional Copy-on-Write (COW) method, with redirect-on-write the actual data activity involved in taking the snapshot is drastically reduced. Moreover, if the size of the data involved in the write request is equal to the system's slot size, there is no need to copy any data at all. If the write request is smaller than the system's slot size, there is still much less copying than with the standard approach of Copy-on-Write.

Figure 1 is an example of the Redirect-on-Write process, the volume is displayed with its data and the pointer to this data.

Figure 1. Redirect-on-Write process: the volume's data and pointer

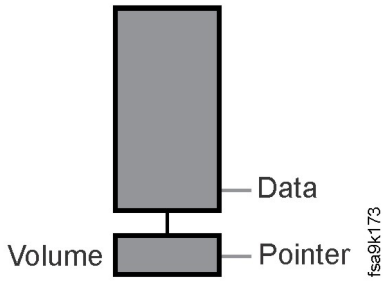

When a snapshot is taken, a new header is written first, as shown in Figure 2.

Figure 2. Redirect-on-Write process: when a snapshot is taken, the header is written first

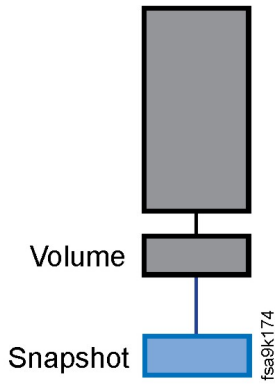

The new data is written anywhere else on the disk, without the need to copy the existing data, as shown in Figure 3.

Figure 3. Redirect-on-Write process: the new data is written

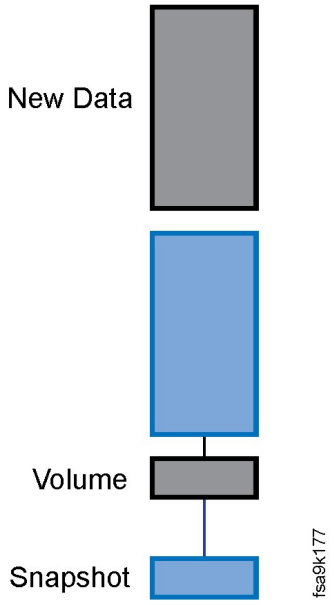

The snapshot points at the old data where the volume points at the new data (the data is regarded as new as it keeps being updated by I/Os), as shown in Figure 4.

Figure 4. Redirect-on-Write process: the snapshot points at the old data where the volume points at the new data

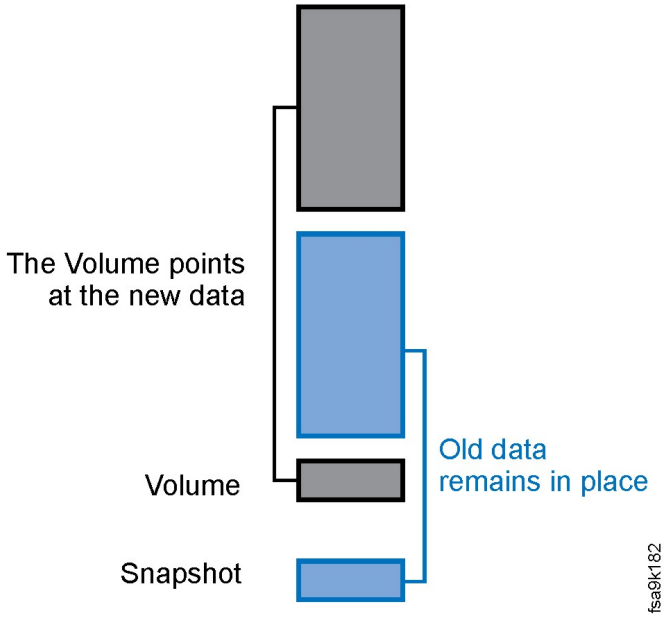

The metadata established at the beginning of the snapshot mechanism is independent of the size of the volume to be copied. This approach allows the user to achieve the following important goals:

Continuous backup
: As snapshots are taken, backup copies of volumes are produced at frequencies that resemble those of *Continuous Data Protection* (CDP). Instant restoration of volumes to virtually any point in time is easily achieved in case of logical data corruption at both the volume level and the file level.

Productivity
: The snapshot mechanism offers an instant and simple method for creating short or long-term copies of a volume for data mining, testing, and external backups.

Full volume copy

Full Volume Copy overwrites an existing volume, and at the time of its creation it is logically equivalent to the source volume.

After the copy is made, both volumes are independent of each other. Hosts can write to either one of them without affecting the other. This is somewhat similar to creating a writable (unlocked) snapshot, with the following differences and similarities:

Creation time and availability
: Both Full Volume Copy and creating a snapshot happen almost instantly. Both the new snapshot and volume are immediately available to the host. This is because at the time of creation, both the source and the destination of the copy operation contain the exact same data and share the same physical storage.

Singularity of the copy operation
: Full Volume Copy is implemented as a single copy operation into an existing volume, overriding its content and potentially its size. The existing target of a volume copy can be mapped to a host. From the host perspective, the content of the volume is changed within a single transaction. In contrast, creating a new writable snapshot creates a new object that has to be mapped to the host.

Space allocation
: With Full Volume Copy, all the required space for the target volume is reserved at the time of the copy. If the storage pool that contains the target volume cannot allocate the required capacity, the operation fails and has no effect. This is unlike writable snapshots, which are different in nature.

Taking snapshots and mirroring the copied volume
: The target of the Full Volume Copy is a Primary volume. This Primary volume can later be used as a source for taking a snapshot or creating a mirror. However, at the time of the copy, neither snapshots nor remote mirrors of the target volume are allowed.

Redirect-on-write implementation
: With both Full Volume Copy and writable snapshots, while one volume is being changed, a redirect-on-write operation will ensure a split so that the other volume maintains the original data.

Performance
: Unlike writable snapshots, with Full Volume Copy, the copying process is performed in the background even if no I/O operations are performed. Within a certain amount of time, the two volumes will use different copies of the data, even though they contain the same logical content. This means that the redirect-on-write overhead of writes occur only before the initial copy is complete. After this initial copy, there is no additional overhead.

Availability
: Full Volume Copy can be performed with source and target volumes in different storage pools.

Restoring volumes and snapshots

The restoration operation provides the user with the ability to instantly recover the data of a Primary volume from any of its locked snapshots.

Restoring volumes

A volume can be restored from any of its snapshots, locked and unlocked. Performing the restoration replicates the selected snapshot onto the volume. As a result of this operation, the Primary volume is an exact replica of the snapshot that restored it.

All other snapshots, old and new, are left unchanged and can be used for further restore operations. A volume can even be restored from a snapshot that has been written to. Figure 1 shows a volume being restored from three different snapshots.

Figure 1. Restoring volumes

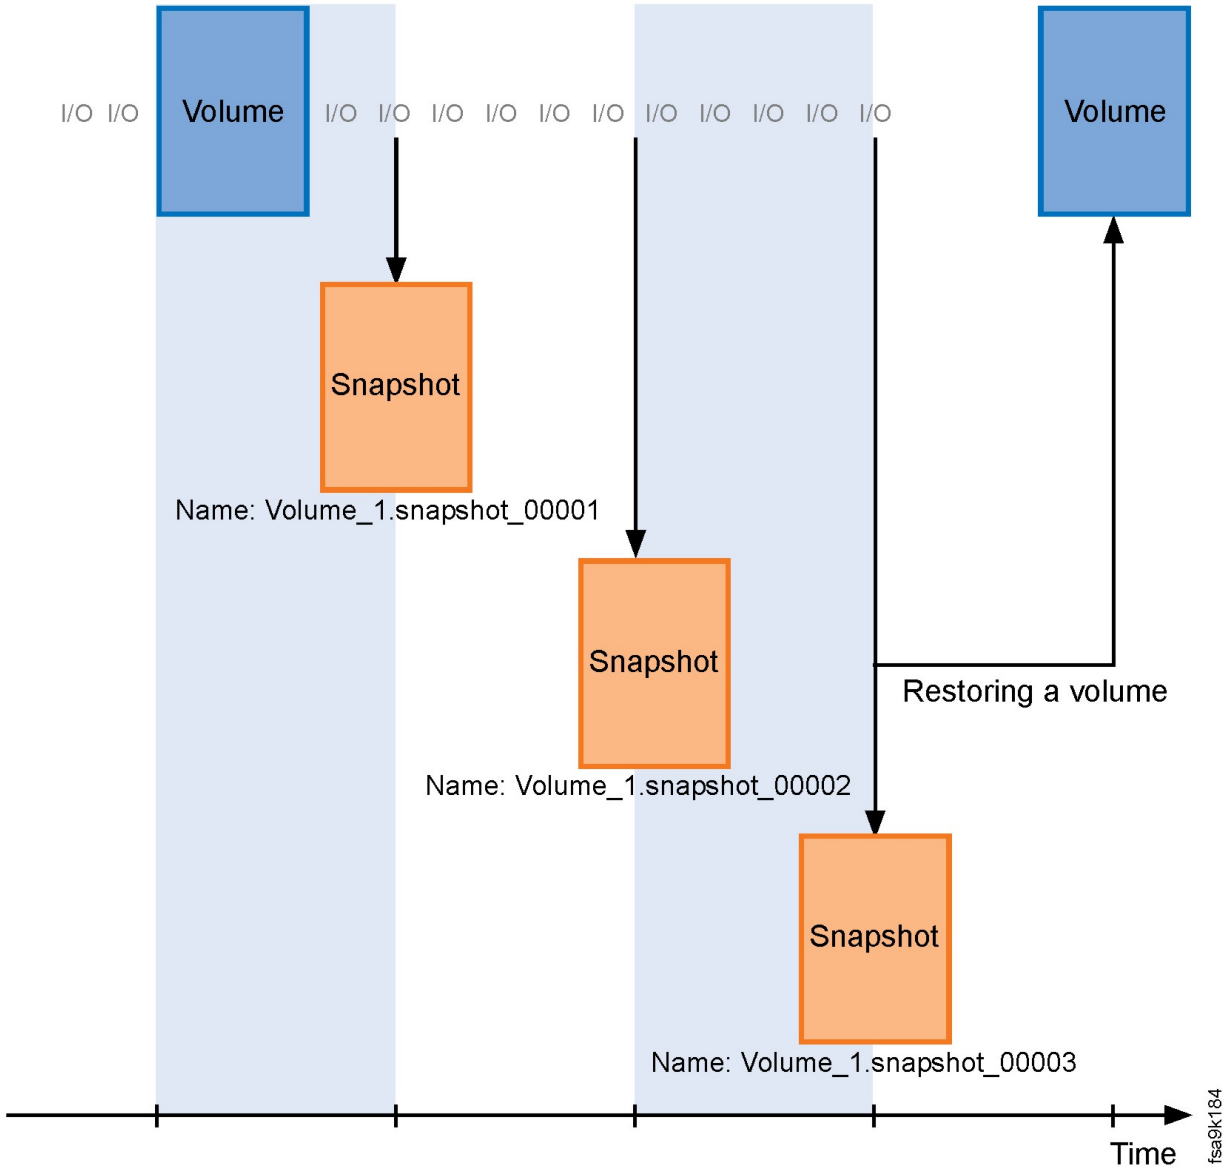

Restoring snapshots

The snapshot itself can also be restored from another snapshot. The restored snapshot retains its name and other attributes. From the host perspective, this restored snapshot is considered an instant replacement of all the snapshot content with other content.

Figure 2 shows a snapshot being restored from two different snapshots.

Figure 2. Restoring snapshots

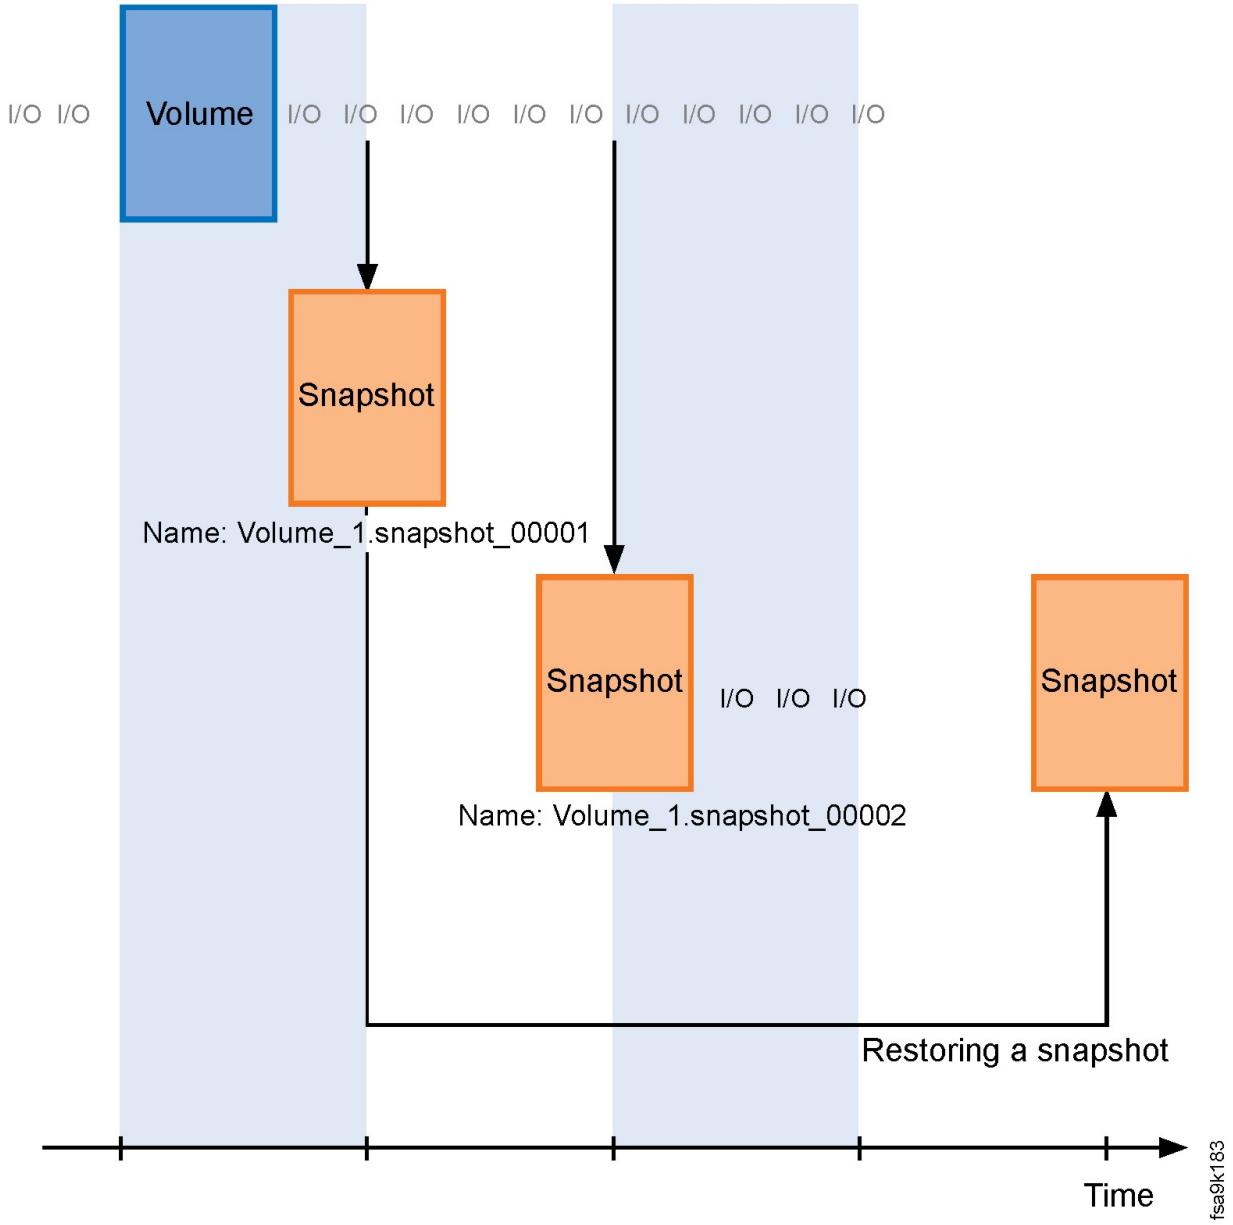

Storage pools

The allocatable capacity of the storage system can be partitioned into *storage pools*, where each volume belongs to a specific storage pool.

The concept of storage pools is purely administrative. Although the hardware resources within the storage system are virtualized in a global sense, the available capacity in the system can be administratively portioned into separate and independent storage pools. Essentially, storage pools function as a virtual way to effectively manage a related group of similarly provisioned logical volumes, consistency groups and their snapshots.

Because of the global deduplication in the storage system (see Flash-optimized data reduction), the capacity assigned to a storage pool is virtual and not totally calculated on the physical or hard capacity of the system. A pool is still configured to be a specific size however the amount of capacity actually used will be dependent on the deduplication and compression savings realized based on the data written from the host system.

Until software version 12.3.0, the upper limit of a storage pool size was 1 TB. From software version 12.3.0 and onwards, the upper limit of the storage pool size is not restricted, provided that the total effective capacity of all pools does not exceed the storage system's maximum effective capacity. However, to protect the system against overallocation of its resources, the maximum allowed total effective capacity of the volumes in a consistency group is limited to 1 PB. The minimum size of a storage pool is 2 TB.

A storage pool can always be grown or shrunk. When decreasing the size of a pool, the only limit is that the new pool size cannot be less than the total size of configured volumes in a pool or not less than the amount of data written for an over-allocated pool.

Storage pools as logical entities

A storage pool is a logical entity and is not associated with a specific flash enclosure or grid controller. All storage pools are equally spread over all flash enclosures and grid controllers in the system.

As a result, there are no limitations on the size of storage pools or on the associations between volumes and storage pools. For example:

- The size of a storage pool can be decreased, limited only by the space consumed by the volumes and snapshots in that storage pool.
- Volumes can be moved between storage pools without any limitations, as long as there is enough free space in the target storage pool.

All of the above transactions are accounting transactions, and do not impose any data copying from one disk drive to another. These transactions are completed instantly.

Moving volumes between storage pools

Volumes that are not in a consistency group (see Consistency groups) can be moved between storage pools without any limitations (assuming adequate space in the new pool). Volumes that are part of a consistency group can be moved together as a group.

For a volume to be moved to a specific storage pool, there must be enough room for it to reside there. If a storage pool is not large enough, the storage pool must be resized, or other volumes must be moved out to make room for the new volume.

A volume and all its snapshots always belong to the same storage pool. Moving a volume between storage pools automatically moves all its snapshots together with the volume.

Protecting snapshots at the storage pool level

Snapshot space must be reserved during pool definition if any of the volumes in the pool will be duplicated with snapshots. This is included as part of the usable capacity in the storage pool. A pool can be resized at a later time to add or remove snapshot space as needed.

Snapshots that participates in the mirroring process can be protected in case of pool space depletion.

This is done by attributing both snapshots (or snapshot groups) and the storage pool with a deletion priority. The snapshots are attributed with a deletion priority between 0 to 4 and the storage pool is configured to disregard snapshots whose priority is above a specific value.

Snapshots with a lower delete priority (higher number) than the configured value might be deleted by the system whenever the pool space depletion mechanism implies so, thus protecting snapshots with a priority equal or higher to this value.

Consistency groups

A *consistency group* is a group of volumes of which a snapshot can be made at the same point in time, therefore ensuring a consistent image of all volumes within the group at that time.

The concept of a consistency group is common among storage systems in which it is necessary to perform concurrent operations collectively across a set of volumes so that the result of the operation preserves the consistency among volumes. For example, effective storage management activities for applications that span multiple volumes, or creating point-in-time backups, is not possible without first employing consistency groups.

The consistency between the volumes in the group is important for maintaining data integrity from the application perspective. By first grouping the application volumes into a consistency group, it is possible to later capture a consistent state of all volumes within that group at a specified point-in-time using a special snapshot command for consistency groups.

Consistency groups can be used to take simultaneous snapshots of multiple volumes, thus ensuring consistent copies of a group of volumes.

A consistency group is also an administrative unit that facilitates simultaneous snapshots of multiple volumes, mirroring of volume groups, and administration of volume sets.

Figure 1 shows the basic workflow of how consistency groups can be created and managed.

Figure 1. Consistency group creation and options

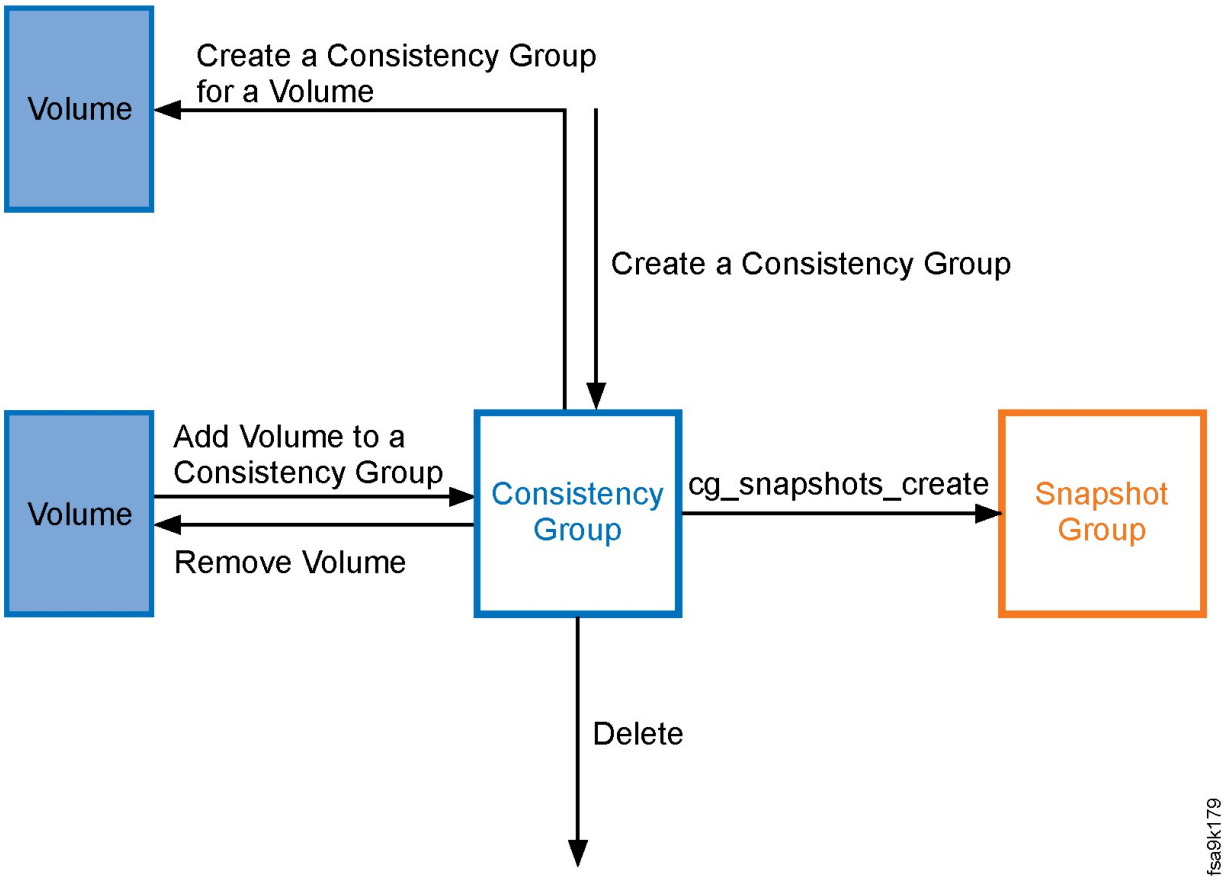

All volumes in a consistency group must belong to the same storage pool. In order to protect the system against overallocation of its resources, the maximum allowed total effective capacity of the volumes in a consistency group is limited to 1 PB (see also Storage pools).

- **Snapshot of a consistency group**
  Taking a snapshot for an entire consistency group means that a snapshot is taken for each volume of the consistency group at the same point-in-time.
- **Consistency group snapshot lifecycle**
  Most snapshot operations can be applied to snapshot groups, where the operation affects every snapshot in the group.

Snapshot of a consistency group

Taking a snapshot for an entire consistency group means that a snapshot is taken for each volume of the consistency group at the same point-in-time.

These snapshots are grouped together to represent the volumes of the consistency group at a specific point in time.

Figure 1. A snapshot is taken for each volume of the consistency group

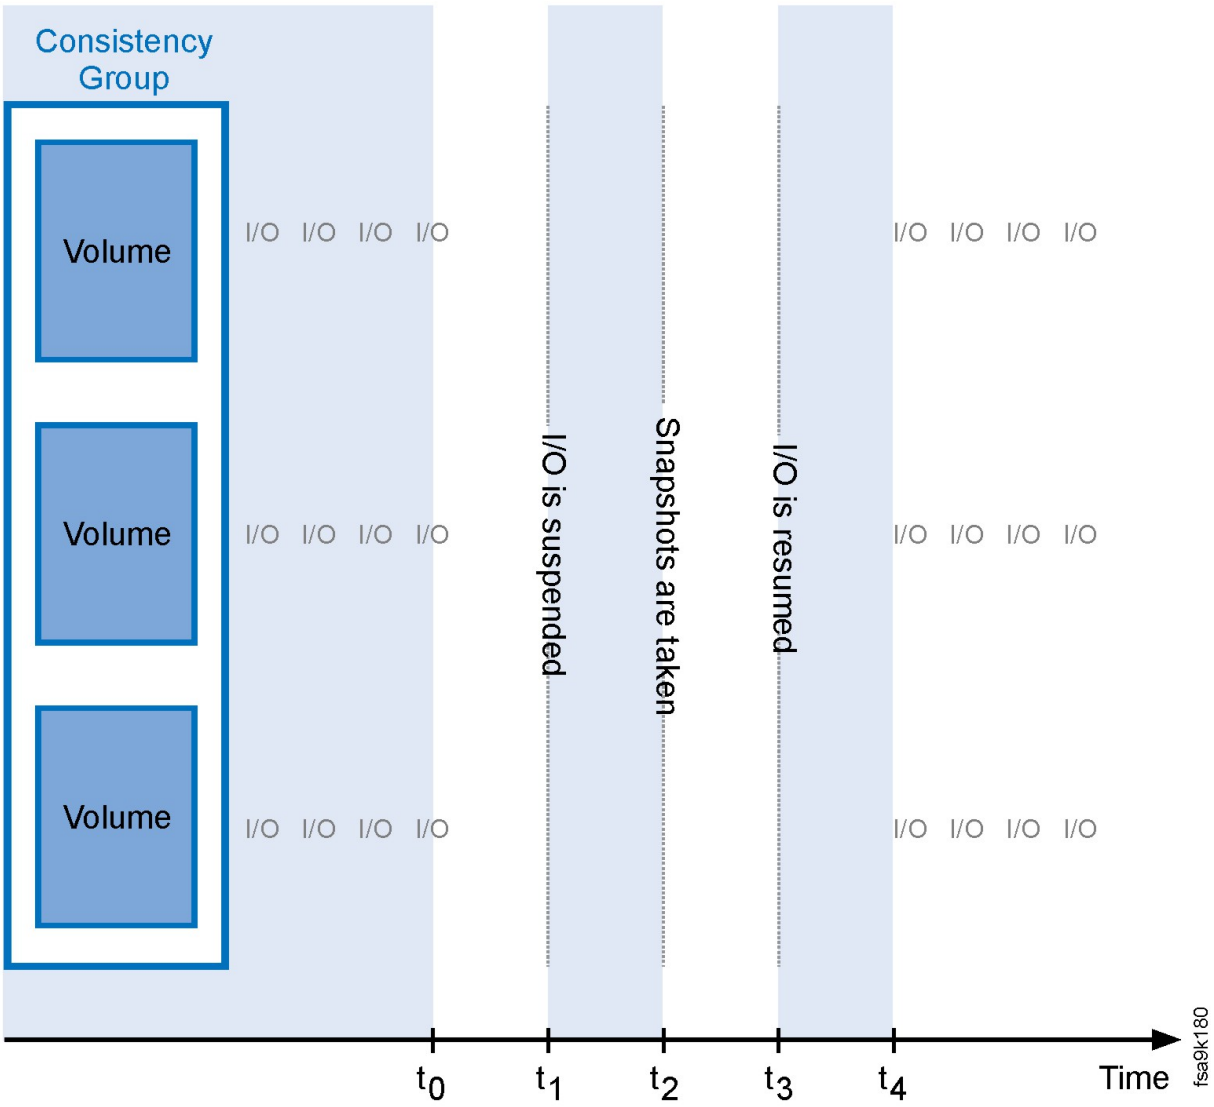

As shown in Figure 1, a snapshot is taken for each of the consistency group's volumes in the following order:

Time = t_0
: Prior to taking the snapshots, all volumes in the consistency group are active and being read from and written to.

Time = t_1
: When the command to snapshot the consistency group is issued, I/O is suspended.

Time = t_2
: Snapshots are taken at the same point in time.

Time = t_3
: I/O is resumed and the volumes continue their normal work.

Time = t_4
: After the snapshots are taken, the volumes resume active state and continue to be read from and written to.

Most snapshot operations can be applied to each snapshot in a grouping, known as a *snapshot set*. The following items are characteristics of a snapshot set:

- A snapshot set can be locked or unlocked. When you lock or unlock a snapshot set, all snapshots in the set are locked or unlocked.
- A snapshot set can be duplicated.
- A snapshot set can be deleted. When a snapshot set is deleted, all snapshots in the set are also deleted.

A snapshot set can be disbanded which makes all the snapshots in the set independent snapshots that can be handled individually. The snapshot set itself is deleted, but the individual snapshots are not.

Consistency group snapshot lifecycle

Most snapshot operations can be applied to snapshot groups, where the operation affects every snapshot in the group.

Figure 1. Most snapshot operations can be applied to snapshot groups

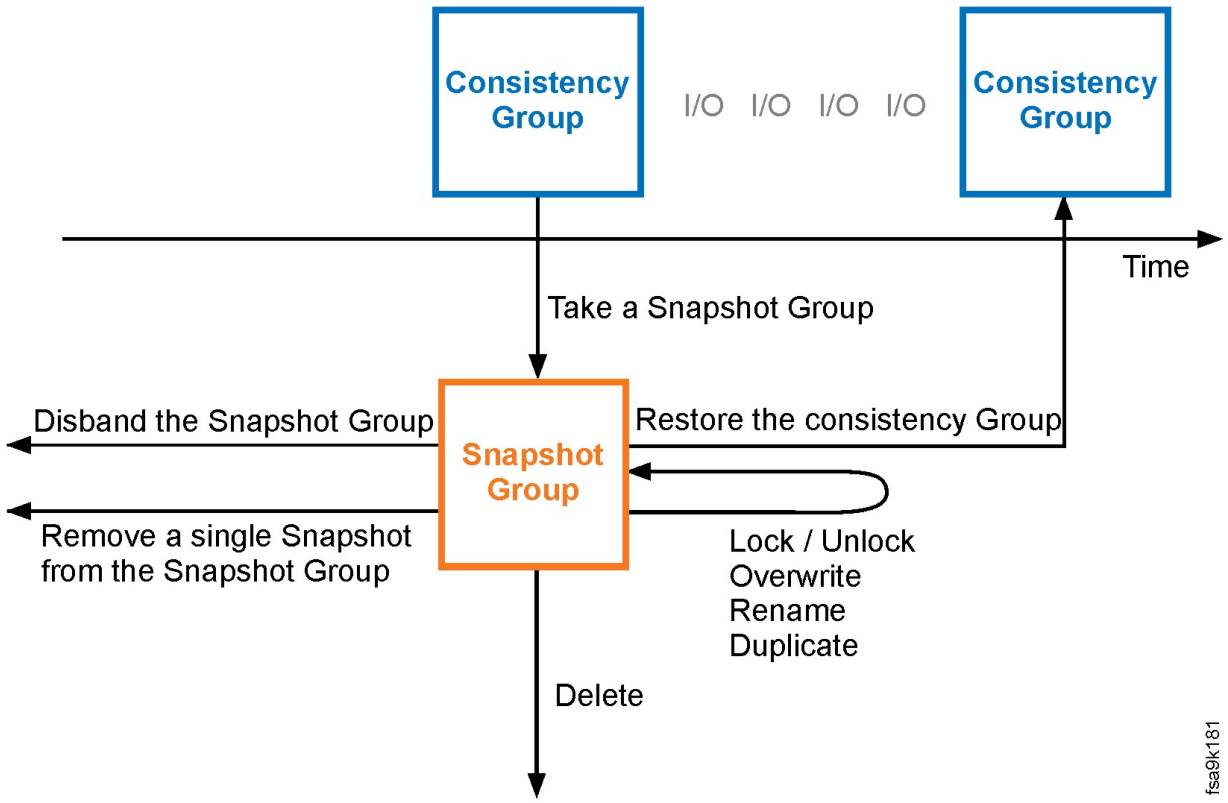

Taking a snapshot group
: Creates a snapshot group.

Restoring consistency group from a snapshot group
: The main purpose of the snapshot group is the ability to restore the entire consistency group at once, ensuring that all volumes are synchronized to the same point in time.
Restoring a consistency group is a single action in which every volume that belongs to the consistency group is restored from a corresponding snapshot that belongs to an associated snapshot group. Not only does the snapshot group have a matching snapshot for each of the volumes, all of the snapshots have the same timestamp. The restored consistency group contains a consistent image of its volumes as they were at a specific point in time.
Note: A consistency group can only be restored from a snapshot group that has a snapshot for each of the volumes. If either the consistency group or the snapshot group has changed after the snapshot group is taken, the restore action does not work.

Listing a snapshot group
: This command lists snapshot groups with their consistency groups and the time the snapshots were taken.
Note: All snapshots within a snapshot group are taken at the same time.

Lock and unlock
: Similar to unlocking and locking an individual snapshot, the snapshot group can be rendered writable, and then be written to. A snapshot group that is unlocked cannot be further used for restoring the consistency group, even if it is locked again.
The snapshot group can be locked again. At this stage, it cannot be used to restore the Primary consistency group. In this situation, the snapshot group functions like a consistency group of its own.

Overwrite
: The snapshot group can be overwritten by another snapshot group.

Rename
: The snapshot group can be renamed.

  Restricted names
  : Do not prefix the snapshot group's name with any of the following strings:

    1. most_recent
    2. last_replicated

Duplicate
: The snapshot group can be duplicated, thus creating another snapshot group for the same consistency group with the time stamp of the first snapshot group.

Disbanding a snapshot group
: The snapshots that comprise the snapshot group are each related to its volume. Although the snapshot group can be rendered inappropriate for restoring the consistency group, the snapshots that comprise it are still attached to their volumes. Disbanding the snapshot group detaches all snapshots from this snapshot group but maintains their individual connections to their volumes. These individual snapshots cannot restore the consistency group, but they can restore its volumes individually.

Changing the snapshot group deletion priority
: Manually sets the deletion priority of the snapshot group.

Deleting the snapshot group
: Deletes the snapshot group along with its snapshots.

Quality of Service (QoS) performance classes

The Quality of Service (QoS) feature allows the storage system to deliver different service levels to hosts that are connected to the same storage system.

The QoS feature favors performance of critical business applications that run concurrently with noncritical applications. Because the flash disk and cache are shared among all applications and all hosts are attached to the same resources, division of these resources among both critical and non-critical applications might have an unintended adverse performance effect on critical applications.

QoS can address this by limiting the rate, based on bandwidth and IOPS, for non-critical applications. Limiting performance resources for non-critical applications means that the remaining resources are available without limitation for the business-critical applications.

The QoS feature is managed through the definition of performance classes and then associating hosts with a performance class. It also applies to storage pools (see Storage pools) and storage domains (see Working with multi-tenancy). Each performance class is now implicitly one of two types: host type or pool/domain type.

The QoS feature possibilities can be summarized as follows:

- Up to 512 performance classes are configurable.
- QoS is applicable to host, domain, pool, volume, and restricted combinations of these entities. For example, hosts cannot be specified for a performance class that already contains a domain or pool.
- Limits can be defined per interface.
- Limits are specified as IOPS or bandwidth.
- Limit calculation is based on preferred practices for setup and zoning.
  The limited I/O processes are expected to always come through all active interface nodes (equal to active interface modules).

Note: For better performance, create one performance class per domain and one performance class per storage pool.

Max bandwidth limit attribute

The host rate limitation group has a max bandwidth limit attribute, which is the number of blocks per second. This number could be either:

- A value between min_rate_limit_bandwidth_blocks_per_sec and max_rate_limit_bandwidth_blocks_per_sec (both are available from the storage system's configuration).
- Zero (0) for unlimited bandwidth.

Connectivity

The storage system connectivity is provided through the physical and other aspects of storage system connectivity.

The following sections provide information about different connectivity aspects, including host and control/management.

- **Physical connectivity**
  The storage system physical connectivity can be provided through the following ports:
- **iSCSI connectivity**
  This section provides a basic explanation of the various iSCSI ports in relation to the available Ethernet ports.
- **Management connectivity**
  The storage system provides multiple ports for management.
- **Virtual LAN (VLAN) support**
  The storage system supports VLANs on Ethernet ports.
- **Host system attachment**
  Hosts of various operating systems can be attached to the storage system over iSCSI or Fibre Channel connections.
- **CHAP authentication of iSCSI hosts**
  Hosts that access the storage system can be authenticated by using the Challenge-Handshake Authentication Protocol (CHAP).
- **Clustering hosts into LUN maps**
  To enhance the management of hosts, the storage system allows clustering hosts together while the clustered hosts are provided with identical mappings.

Physical connectivity

The storage system physical connectivity can be provided through the following ports:

- 16-Gigabit (16 Gbps) Fibre Channel (FC) ports for host-based I/O using the FC protocol (FCP) over Fibre Channel networks
- 10-Gigabit Ethernet (10 Gbps) ports for host-based I/O using the iSCSI protocol over Ethernet networks
- 1-Gigabit Ethernet for management (Hyper-Scale Manager UI or CLI) connectivity (for more information, see Management and monitoring)
- 1-Gigabit Ethernet port for incoming-only traffic of technician maintenance operations
- Two 1-Gigabit Ethernet ports for remote support (for more information, see Remote support and proactive support)

The total number and type of host connectivity ports depends on the number of grid controllers in the storage system configuration. Table 1 breaks down each model, feature code, and number of connectivity ports.

Table 1. Number of connectivity ports

Model	Feature code	Grid controller port configuration	Number of grid controllers

Model	Feature code	Grid controller port configuration	Number of grid controllers
FlashSystem A9000 model 415	5010	4 Fibre Channel, 2 iSCSI	3
	5011	4 iSCSI	3
FlashSystem A9000 models 425/U25	5003	4 Fibre Channel, 2 iSCSI	3
	5004	4 iSCSI	3
	5005	4 NVMe/FC, 2 iSCSI	3
FlashSystem A9000R model 415	5001	4 Fibre Channel, 2 iSCSI	4, 6, 8, or 12
	5002	4 iSCSI	4, 6, 8, or 12
FlashSystem A9000R 425/U25	5003	4 Fibre Channel, 2 iSCSI	3, 4, 6, or 8
	5004	4 iSCSI	3, 4, 6, or 8
	5005	4 NVMe/FC, 2 iSCSI	3, 4, 6, or 8

IBM FlashSystem® A9000R storage systems are shipped with grid controllers that feature NVMe/FC ready adapters (see the Fibre Channel NVMe connectivity hardware announcement). NVM Express® (NVMe) allows servers to leverage the native parallelism of flash media, reduces overall I/O overhead, and increases bandwidth. NVMe/FC enables NVMe over a Fibre Channel (FC) network fabric, thus combining the benefits of all-flash SAN storage with NVMe performance over existing infrastructure.
Note: To support provisioning of NVMe/FC-based connectivity , a future software upgrade will be required.

iSCSI connectivity

This section provides a basic explanation of the various iSCSI ports in relation to the available Ethernet ports.

The storage system uses 10-Gigabit Ethernet ports for iSCSI and other connectivity uses, like management. The number of iSCSI ports in a fully equipped storage system depends on the grid controller port configuration. Table 1 breaks down each model, feature code, and maximum number of Ethernet ports.

Table 1. Ethernet ports for iSCSI service

Model	Feature code	Grid controller port configuration	Maximum number of Ethernet ports
FlashSystem A9000 model 415	5010	4 Fibre Channel, 2 iSCSI	6
	5011	4 iSCSI	12
FlashSystem A9000 models 425	5003	4 Fibre Channel, 2 iSCSI	6
	5004	4 iSCSI	12
	5005	4 NVMe/FC, 2 iSCSI	6
FlashSystem A9000R model 415	5001	4 Fibre Channel, 2 iSCSI	24
	5002	4 iSCSI	48
FlashSystem A9000R 425/U25	5003	4 Fibre Channel, 2 iSCSI	16
	5004	4 iSCSI	32
	5005	4 NVMe/FC, 2 iSCSI	16

There ports should connect to the user's network and provide connectivity to the iSCSI hosts. The iSCSI ports can also accept management connections.

Management connectivity

The storage system provides multiple ports for management.

Management connectivity is used for the following functions, including:

- Issuing CLI commands through the XCLI utility.
- Controlling the storage system through the IBM® Hyper-Scale Manager UI.
- Sending e-mail and SMS notification messages and SNMP traps about event alerts.
- Connecting with NTP, Key Server, DNS, and LDAP.

Management connectivity uses these ports:

Gigabit Ethernet management ports
: These ports are dedicated for CLI and Hyper-Scale Manager UI communications, as well as for outgoing SNMP and SMTP connections. The FlashSystem A9000 has two management ports, and the FlashSystem A9000R has three management ports.

Field technician ports
: These Ethernet ports are used for incoming management traffic only. The ports are utilized only for the field technician's laptop computer, and must not be connected to the user's IP network.

Remote support ports
: These two ports are used for remote support by IBM technical support.

To ensure management redundancy in case of module failure, the storage system management function is accessible from two different IP addresses in IBM FlashSystem® A9000, and three in IBM FlashSystem A9000R. Each of the three IP addresses is handled by a different hardware module. The various IP addresses are transparent to the user and management functions can be performed through any of the IP addresses. These addresses can be accessed simultaneously by multiple clients. Users only need to configure the set of management IP addresses that are defined for the specific system.
Note: All management IP interfaces must be connected to the same subnet and use the same network mask, gateway, and MTU.
The management connectivity allows users to manage the system from both the CLI and GUI. Accordingly, both can be configured to manage the system through iSCSI IP interfaces. Both CLI and GUI management is run over TCP port 7778. With all traffic encrypted through the Secure Sockets Layer (SSL) protocol.

Optional IPv6 connectivity for management traffic

The storage system supports optional IPv6 through stateless auto-configuration and full IPsec (IKE2, transport, and tunnel mode) for the management ports. IPv6 is not supported for the iSCSI or field technician ports.

When IPv6 is enabled, stateless auto-configuration is automatically enabled as well, and the system interfaces are getting ready to work with IPv6. Accordingly, when looking for DNS addresses, the system also looks for AAAA entries. In addition, each IP interface in the system may have several IP addresses: static IPv4 address, static IPv6 address, and the stateless configuration link and site local IPv6 addresses. Where multiple IPv6 static addresses are assigned for each interface, the system supports only one address per interface.

Programs that are using connections on the management and VPN ports must support IPv6 addresses.

System-initiated IP communication

The storage system can also initiate IP communications to send event alerts as necessary, for example call-home. Two types of system-initiated IP communications exist:

Sending e-mail notifications through the SMTP protocol
: E-mails are used for both e-mail notifications and for SMS notifications through the SMTP to SMS gateways.

Sending SNMP traps
: SMTP and SNMP communications can be initiated from any of the three IP addresses. This is different from the CLI and GUI, which are user initiated. Accordingly, it is important to configure all three IP interfaces and to verify that they have network connectivity.

Virtual LAN (VLAN) support

The storage system supports VLANs on Ethernet ports.

VLANs in the storage system enable concurrent connectivity of multiple VLANs to one storage system. This is achieved by allowing multiple virtual IP addresses per physical Ethernet port (otherwise known as Port Trunking), and routing IP packets to and from the virtual addresses based on their VLAN tags. This capability is enabled by default.

Table 1 shows the VLAN feature capacities.

Table 1. FlashSystem A9000R and A9000R VLAN support specifications

VLAN feature	Capability	
Maximum number of VLANs	100	
Maximum number of IP interfaces	512	
Maximum number of IP interfaces per Ethernet port	Host connectivity	100
	Replication	1
	Non-data ports	1
Maximum number of IP interfaces per VLAN	Per system	12
	Per port	1
Default VLAN priority on iSCSI ports	6	
VLAN number range	1-4094	

Note: Multiple IPs cannot be assigned to the same VLAN and Ethernet port. For example, if 192.0.2.1 is assigned to Ethernet port 3 and VLAN 100, then 192.0.2.2 cannot be assigned to Ethernet port 3 and VLAN 100, but can be assigned to either a different Ethernet port on VLAN 100, or a different VLAN on Ethernet port 3.

Host system attachment

Hosts of various operating systems can be attached to the storage system over iSCSI or Fibre Channel connections.

The following subtopics provide information about different host system attachment aspects.

Balanced traffic and no single point of failure

Although the storage system distributes the traffic across all system modules, the storage administrator is responsible for ensuring that host I/O operations are equally distributed among the different interface modules.

The workload balance should be monitored and reviewed when host traffic patterns change. The storage system does not automatically balance incoming host traffic.

The storage administrator is responsible for ensuring that host connections are made redundantly in such a way that a single failure, such as in a module or HBA, will not cause all paths to the machine to fail. In addition, the storage administrator is responsible for making sure that the host workload is adequately spread across the different connections and interface modules.

Attaching volumes to hosts

While the storage system identifies volumes and snapshots by name (see Volumes and snapshots), hosts identify volumes and snapshots according to their logical unit number (LUN).

A *LUN* is an integer that is used when attaching a system's volume to a registered host. Each host can access some or all of the volumes and snapshots on the storage system, up to a set maximum. Each accessed volume or snapshot is identified by the host through a LUN.

For each host, a LUN identifies a single volume or snapshot. However, different hosts can use the same LUN to access different volumes or snapshots.

Use your operating system's native host attachment for volume attachment. When using host attachment, be sure to apply the multipath connectivity configurations on the host. For more information, see Host network connectivity and multipathing configurations and Configuring host attachments.
Important: LUN0 can be mapped to a volume like any other LUN however it is better to not use LUN0 as a normal LUN.

Multipathing

The system's host connectivity interfaces utilize multipathing access algorithms. When a host connects to the system through several independent ports, each volume can be accessed directly through any of the host connectivity interfaces, and no further interaction is required. This allows using more than one physical path for transferring data between the host and the storage system.

CHAP authentication of iSCSI hosts

Hosts that access the storage system can be authenticated by using the Challenge-Handshake Authentication Protocol (CHAP).

When CHAP support is enabled, hosts are securely authenticated by the storage system. This increases overall system security by verifying that only authenticated parties are involved in host-storage interactions.

Definitions

CHAP authentication
: An authentication process of an iSCSI initiator by a target through comparing a secret hash that the initiator submits with a computed hash of that initiator's secret which is stored on the target.

Initiator
: The host.

One-way (unidirectional CHAP)
: CHAP authentication where initiators are authenticated by the target, but not vice versa.

Supported configurations

CHAP authentication type
: One-way (unidirectional) authentication mode, meaning that the Initiator (host) has to be authenticated by the storage system.

MD5
: CHAP authentication utilizes the MD5 hashing algorithm.

Access scope
: CHAP-authenticated initiators are granted access to the storage system by defining mapping that may restrict access to some volumes.

Authentication modes

The following authentication modes are supported:

None (default)
: In this mode, an initiator is not authenticated by the storage system.

CHAP (one-way)
: In this mode, an initiator is authenticated by the storage system based on the pertinent initiator's submitted hash, which is compared to the hash computed from the initiator's secret stored on the storage system.

Changing the authentication mode from None to CHAP requires an authentication of the host. Changing the mode from CHAP to None does not require an authentication.

Complying with RFC 3720

The CHAP authentication procedure complies with the CHAP requirements as defined in IETF RFC 3720.

Secret length
: The secret key has to be between 96 bits and 128 bits.

Initiator secret uniqueness
: Upon defining or updating an initiator (host) secret, the system compares the entered secret's hash with existing secrets stored by the system and determines whether the secret is unique. If it is not unique, the system presents a warning to the user, but does not prevent the command from completing successfully.

Clustering hosts into LUN maps

To enhance the management of hosts, the storage system allows clustering hosts together while the clustered hosts are provided with identical mappings.

The mapping of volumes to LUN identifiers is defined per cluster and applies to all of the hosts in the cluster.

Figure 1 shows how hosts are added to and removed from a cluster.

Figure 1. A volume, a LUN, and cluster hosts

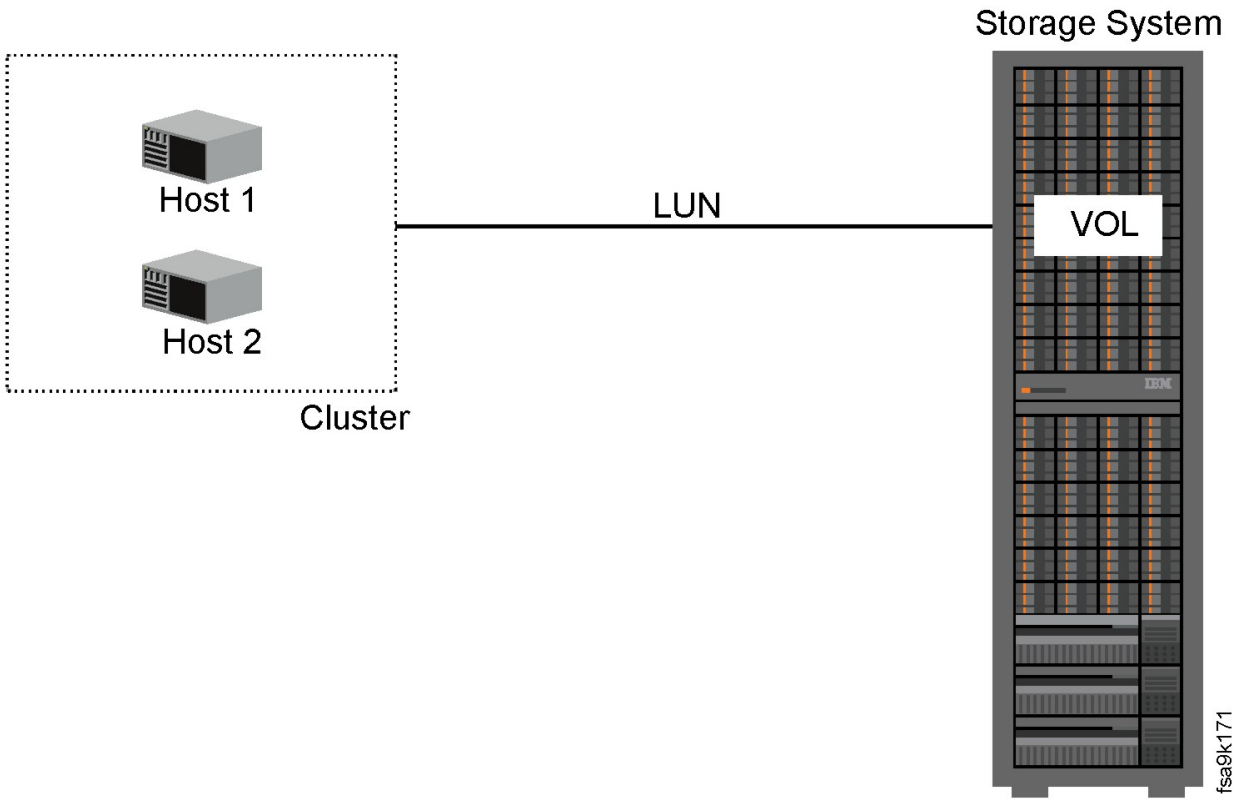

Adding a host to a cluster

Adding a host to a cluster is a straightforward action in which a host is added to a cluster and is connected to a LUN. This results in either:

- Changing the host mapping to the cluster mapping.
- Changing the cluster mapping to be identical to the mapping of the newly added host.

Removing a host from a cluster

When a host is removed from the cluster:

- The removed host mapping remains identical to the LUN mapping of the cluster.
- The host mapping definitions do not revert to the original mapping (the mapping that was in effect before the host was added to the cluster).
- The host LUN mapping can be changed.

The following restrictions are applicable:

- The storage system defines the same mapping to all of the hosts of the same cluster. No hierarchy of clusters is maintained.
- As shown inFigure 2, a volume cannot be mapped to an already mapped LUN.

Figure 2. Volume that cannot be mapped to an already mapped LUN

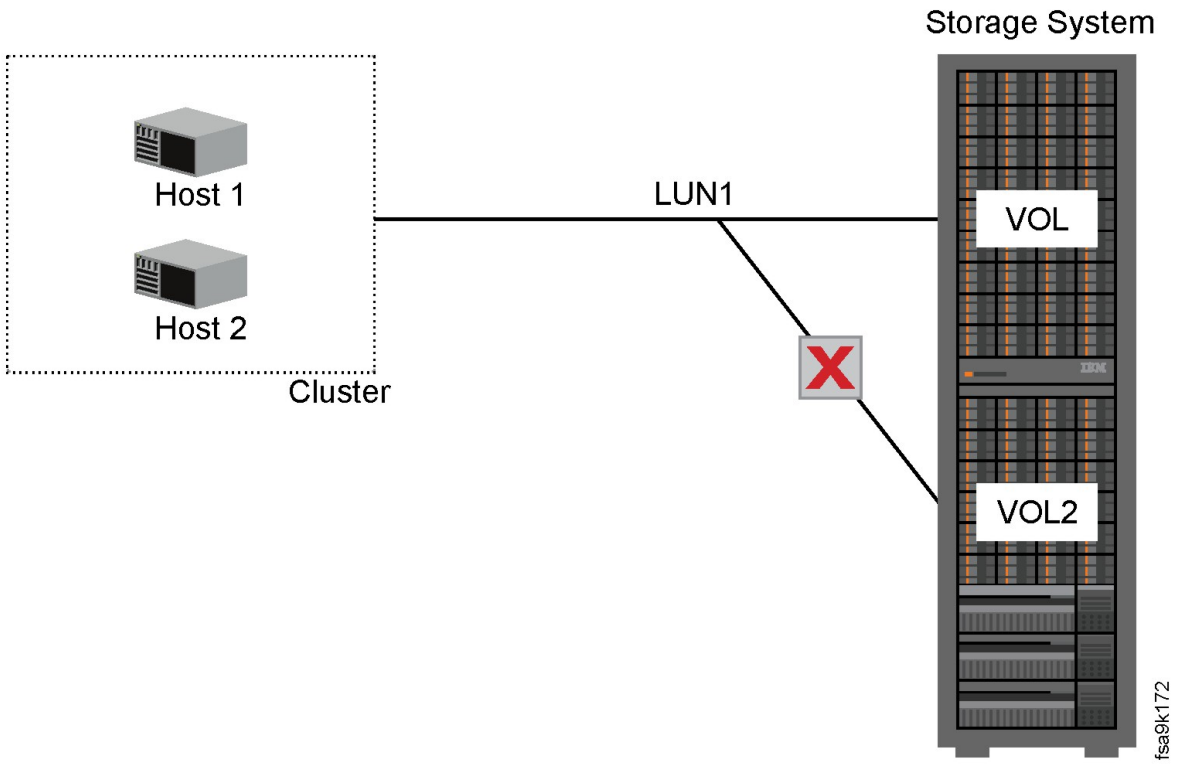

- As shown in Figure 3, a mapped volume cannot be mapped to another LUN.

Figure 3. Mapped volume that cannot be mapped to another LUN

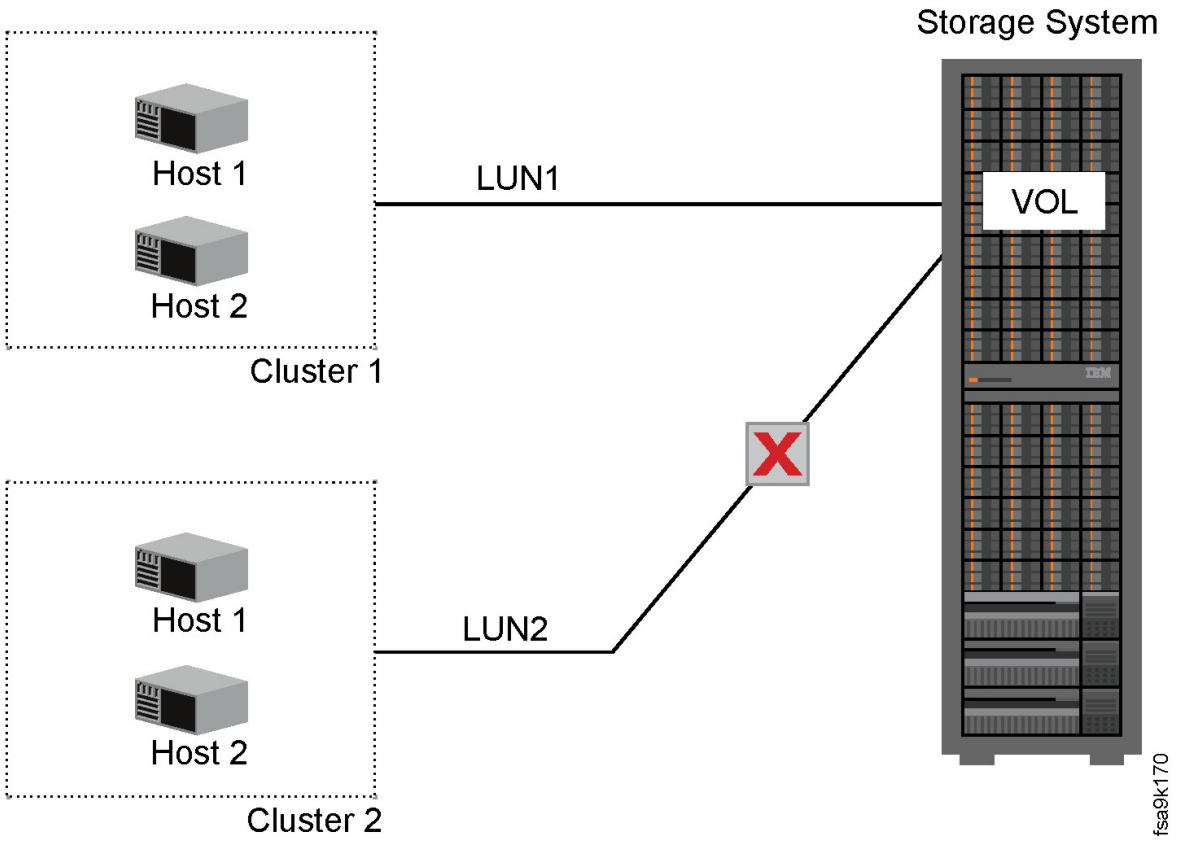

Synchronous remote mirroring

Remote mirroring allows replication of data between two geographically remote sites, allowing full data recovery from the remote site in different disaster scenarios.

The process of ensuring that both storage systems contain identical data at all times is called *remote mirroring*. Remote mirroring can be established between two remote storage systems to provide data protection for the following types of site disasters:

Local site failure
: When a disaster occurs at a certain site, the remote site takes over and maintains full service to the hosts connected to the original site. The mirroring is resumed after the failing site recovers.

Split-brain scenario
: After a communication loss between the two sites, each site maintains full service to the hosts. After the connection is resumed and the link (mirror) is established, the sites complement each other's data to regain full synchronization.

For the information on the cross-system use of synchronous remote mirroring, see the feature availability matrix.

- **Remote mirroring basic concepts**
  Synchronous remote mirroring provides continuous availability of critical information in the case of a disaster scenario.
- **Synchronous mirroring operations**
  Remote mirroring operations involve *configuration*, *initialization*, *ongoing operation*, *handling of communication failures*, and *role switching*.
- **Synchronous mirroring configuration and activation options**
  The remote mirroring configuration process involves configuring volumes and volume pairs.
- **Synchronous mirroring statuses**
  The status of a synchronous remote mirroring volume depends on the communication link and on the coupling between the Primary volume and the Secondary volume.
- **Synchronous mirroring role switchover and role change**
  When role switchover occurs, the Primary volume becomes the Secondary volume, and the Secondary volume becomes the Primary volume.
- **I/O operations in synchronous mirroring**
  I/O operations are performed on the Primary and Secondary volumes across various configuration options.
- **Coupling synchronization process**
  When a failure condition has been resolved, remote mirroring begins the process of synchronizing the coupling. This process updates the Secondary volume with all the changes that occurred while the coupling was not operational.
- **Synchronous mirroring of consistency groups**
  Mirroring can be applied to whole consistency groups.

Remote mirroring basic concepts

Synchronous remote mirroring provides continuous availability of critical information in the case of a disaster scenario.

A typical remote mirroring configuration involves the following two sites:

Primary site
: The location of the primary storage system.
A local site that contains both the data and the active servers.

Servers may simultaneously perform primary or secondary roles with respect to their hosts. As a result, a server at one site can be the primary storage system for a specific application, while simultaneously being the secondary storage system for another application.

Secondary site
: The location of the secondary backup storage system.
A remote site that contains a copy of the data and standby servers. Following a disaster at the primary site, the servers at the secondary site become active and start using the copy of the data.

Primary volume
: The volume which is mirrored. The Primary volume is usually located at the primary site.

Secondary volume
: The volume to which the Primary volume is mirrored. The Secondary volume is usually located at the secondary site.

Synchronous remote mirroring is performed during each write operation. The write operation issued by a host is applied to both the primary and the secondary storage systems. Figure 1 shows a synchronous remote mirroring scheme.

Figure 1. Synchronous remote mirroring scheme

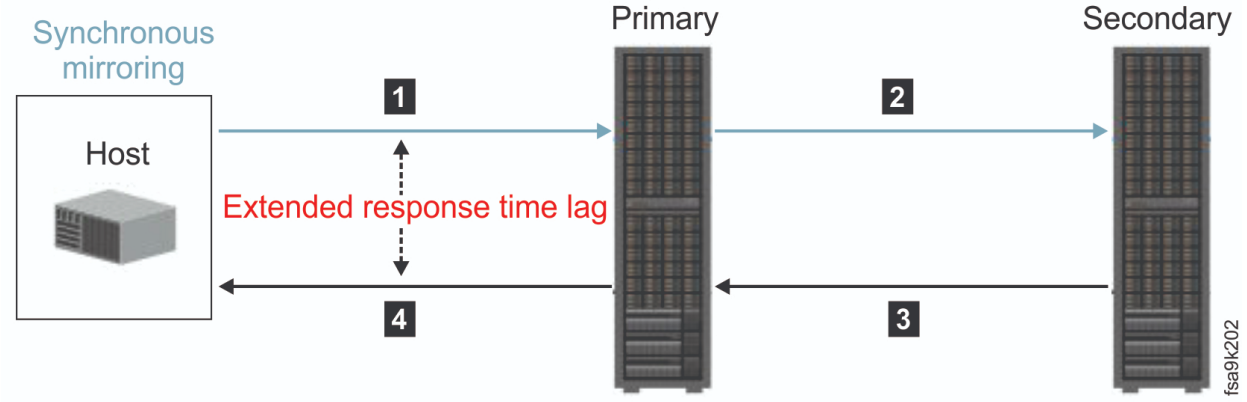

Note: When using remote mirroring with FlashSystem A9000 or A9000R, data is transferred over the mirror connectivity in uncompressed format. The data is deduplicated and compressed again after it reaches the remote system.

When a volume is mirrored, reading is performed from the Primary volume, while writing is performed on both the Primary and the Secondary volumes, as previously described.

Synchronous mirroring operations

Remote mirroring operations involve *configuration*, *initialization*, *ongoing operation*, *handling of communication failures*, and *role switching*.

The following list describes the remote mirroring operations:

Configuration
: Configuration is the act of defining Primary and Secondary volumes for a mirror relation.

Initialization
: Remote mirroring operations begin with a Primary volume that contains data and a new Secondary volume. Next, data is copied from the Primary volume to the Secondary volume. This process is called *initialization*. Initialization is performed once in the lifetime of a remote mirroring coupling. After it is successfully completed, both volumes are synchronized.

Ongoing operation
: After the *initialization* process is complete, remote mirroring is activated. During this activity, all data is written to the Primary volume and to the Secondary volume. The write operation is complete after an acknowledgment is received from the Secondary volume. At any point, the Primary and Secondary volumes contain identical data except for any unacknowledged (pending) writes.

Handling of communication failures
: Communication between sites may break. In this case, the primary site continues its function and updates the secondary site after communication resumes. This process is called *synchronization*.

Role switching
: When needed, a volume can change its role from Primary to Secondary or vice versa, either as a result of a disaster at the primary site, maintenance operations, or intentionally, to test the disaster recovery procedures.

Using snapshots in synchronous mirroring

The storage system uses snapshots to identify inconsistencies that may arise between updates.

If the link between volumes is disrupted or if the mirroring is deactivated, the Primary volume continues accepting host writes, but does not replicate the writes onto the Secondary volume. After the mirroring is restored and activated, the system takes a snapshot of the Secondary volume, which represents the data that is known to be mirrored. This snapshot is called the *last-consistent snapshot*. Only then more recent writes to the Primary volume are replicated to the Secondary volume.

The last-consistent snapshot is automatically deleted after the resynchronization is complete for all mirrors on the same target. However, if the Secondary volume role is changed to Primary during resynchronization, the last-consistent snapshot will not be deleted.

Synchronous mirroring configuration and activation options

The remote mirroring configuration process involves configuring volumes and volume pairs.

Volume configuration

The following concepts are to be configured for volumes and the relations between them:

The *volume role* is the current function of the volume. The following volume roles are available:

None
: The volume is created using normal volume creation procedures and is not mirrored.

Primary
: The volume is directly written to by the host.

Secondary
: A backup to the Primary volume.
Data can be read from the Secondary volume by a host. Data cannot be written to the Secondary volume by any host.

Mixed configuration
: In some cases, the volumes on a single storage system can be defined in a *mixed configuration*. For example, a storage system can contain volumes whose role is defined as Primary, as well as volumes whose role is defined as Secondary. In addition, some volumes might not be involved in a remote mirroring coupling at all.

Configuration error
: In some cases, configuration on both sides might be changed in a non-compatible way. This is defined as a *configuration error*. For example, switching the role of only one side when communication is down causes a configuration error when connection resumes, because each side is configured as a Primary or Secondary.

Coupling activation

When a pair of volumes point to each other, it is referred to as a *coupling*. In a *coupling relation*, two volumes, referred to as *peers*, participate in a remote mirroring system with the Secondary peer serving as the backup for the Primary peer. The coupling configuration is identical for both Primary and Secondary volumes.

Remote mirroring can be manually activated and deactivated per coupling. When activated, the coupling is in *Active* mode. When deactivated, the coupling is in *Standby* mode.

These modes have the following functions:

Active
: Remote mirroring is functioning and the data is replicated.

Standby
: Remote mirroring is deactivated. The data is not replicated to the Secondary volume.
Standby mode is used mainly when maintenance is performed on the secondary site or during communication failures between the sites. In this mode, the Primary volumes will not generate mirroring-failure alerts.

The coupling lifecycle has the following characteristics:

- When a coupling is created, it is always initially in *Standby* mode.
- Only a coupling in *Standby* mode can be deleted.

Supported network configurations

Synchronous mirroring supports the following network configurations:

- Either Fibre Channel (FC) or iSCSI connectivity can be used for replication, regardless of the connectivity that is used by the host to access the Primary volume.
- The remote system must be defined in the *remote target connectivity* definitions.
- All the volumes that belong to the same consistency group must reside on the same remote system.
- Primary and Secondary volumes must have exactly the same size.

Synchronous mirroring statuses

The status of a synchronous remote mirroring volume depends on the communication link and on the coupling between the Primary volume and the Secondary volume.

Table 1 lists the different statuses of a synchronous remote mirroring volume during remote mirroring operations.

Table 1. Synchronous mirroring statuses

Entity	Status type	Possible status values	Description
Link	Operational status	• Up • Down	Specifies if the communications link is up or down. The link status of the Primary volume is also the link status of the Secondary volume.
Coupling	Operational status	• Operational • Non-operational	Specifies if remote mirroring is working. To be operational, the link status must be up and the coupling must be activated. If the link is down or if the remote mirroring feature is in Standby mode, the status is Non-operational.
	Synchronization status	• Initialization • Synchronized • Unsynchronized • Consistent • Inconsistent	For detailed description of each status, see "Synchronization status" below.
	Last-secondary timestamp	Point-in-time date	Timestamp for when the secondary volume was last synchronized.
	Synchronization progress	Synchronization status	The relative portion of data remaining to be synchronized between the Primary and Secondary volumes due to non-operational coupling.
	Secondary-locked	Boolean	If the Secondary volume is locked for writing due to lack of space, the *Secondary-locked* status is *true*. This may occur during the synchronization process, when there is not enough space for the last-consistent snapshot. Otherwise, the *Secondary-locked* status is *false*.
	Configuration error	Boolean	If the configuration of the Primary and Secondary volumes is inconsistent., the *Configuration error* status is *true*.

Synchronization status

The synchronization status reflects the consistency of the data between the Primary and Secondary volumes.

Because remote mirroring is for ensuring that the Secondary volume is an identical copy of the Primary volume, this status indicates whether this objective is currently attained.

The possible synchronization statuses for the Primary volume are:

Initialization
: The first step in remote mirroring is to create a copy of the data from the Primary volume to the Secondary volume. During this step, the coupling status remains *Initialization*.

Synchronized (Primary volume only)
: This status indicates that all data that was written to the Primary volume and acknowledged has also been written to the Secondary volume. Ideally, the Primary and Secondary volumes should always be synchronized. This does not imply that the two volumes are identical because at any time there might be a limited amount of data that was written to one volume, but was not yet acknowledged by the Secondary volume. These are also known as *pending writes*.

Unsynchronized (Primary volume only)
: After a volume has completed the *Initialization* stage and achieved the *Synchronized* status, a volume can become unsynchronized. This occurs when it is not known whether all the data that was written to the Primary volume was also written to the Secondary volume. This status occurs in the following cases:

  - **Communications link is down** – As a result of the communication link going down, some data might have been written to the Primary volume, but was not yet replicated to the Secondary volume.
  - **Secondary system is down** – This is similar to communication link errors because in this state, the primary system is updated while the secondary system is not.
  - **Remote mirroring is deactivated** – As a result of the remote mirroring deactivation, some data might have been written to the Primary volume and not to the Secondary volume.

Consistent
: The Secondary volume is an identical copy of the Primary volume.

Inconsistent
: There is a discrepancy between the data on the Primary and Secondary volumes.

It is always possible to reestablish the Synchronized status when the link is reestablished or the remote mirroring feature is reactivated, no matter what was the reason for the *Unsynchronized* status.

Because all updates to the Primary volume that are not written to the Secondary volume are recorded, these updates are written to the Secondary volume. The synchronization status remains *Unsynchronized* from the time that the coupling is not operational until the synchronization process is completed successfully.

Last-secondary timestamp

A timestamp is taken when the coupling between the Primary and Secondary volumes becomes non-operational.

This timestamp specifies the last time that the Secondary volume was consistent with the Primary volume. This status has no meaning if the coupling's synchronization state is still *Initialization*.

For synchronized coupling, this timestamp specifies the current time. Most importantly, for an unsynchronized coupling, this timestamp denotes the time when the coupling became non-operational.

The timestamp is returned to current only after the coupling is operational and the Primary and Secondary volumes are synchronized.

Synchronization progress

During the synchronization process, when the Secondary volumes are being updated with previously written data, the volumes are given a dynamic synchronization process status.

This status comprises the following sub-statuses:

Size to complete
: The size of data that requires synchronization.

Part to synchronize
: The size to synchronize divided by the maximum size-to-synchronize since the last time the synchronization process started. For coupling initialization, the size-to-synchronize is divided by the volume size.

Time to synchronize
: Time estimation that is required to complete the synchronization process and achieve synchronization, based on past rate.

Secondary-locked error status

When synchronization is in progress, there is a period in which the Secondary volume is not consistent with the Primary volume. While in this state, the Secondary volume maintains a last-consistent snapshot. Provided that every I/O operation requires a copy-on-write partition, this may result in insufficient space and, consequently, in the failure of I/O operations to the Secondary volume.

Whenever I/O operations to the Secondary volume fail due to insufficient space, all couplings in the system are set to the *Secondary-locked* status and become non-operational. The administrator is notified of a critical event, and can free space on the system containing the Secondary volume.

Synchronous mirroring role switchover and role change

When role switchover occurs, the Primary volume becomes the Secondary volume, and the Secondary volume becomes the Primary volume.

Role switching can occur when the synchronous remote mirroring function is either operational or not operational, as described in the following sections.

- **Role switchover when synchronous mirroring is operational**
  When the remote mirroring function is operational, role switching between Primary and Secondary volumes can be initiated from the management IBM Hyper-Scale Manager or CLI.

- **Role change when synchronous mirroring is not operational**
  A more complex situation for role switching is when there is no communication between the two sites, either because of a network malfunction, or because the primary site is no longer operational.

Role switchover when synchronous mirroring is operational

When the remote mirroring function is operational, role switching between Primary and Secondary volumes can be initiated from the management IBM Hyper-Scale Manager or CLI.

There are two typical reasons for performing a switchover when communication between the volumes exists:

Drills
: Drills can be performed on a regular basis to test the functioning of the secondary site. In a drill, an administrator simulates a disaster and tests that all procedures are operating smoothly.

Scheduled maintenance
: To perform maintenance at the primary site, switch operations to the secondary site on the day before the maintenance. This can be done as a preemptive measure when a primary site problem is known to occur.

The CLI command that performs the role switchover must be run on the Primary volume. The switchover cannot be performed if the Primary and Secondary volumes are not synchronized.

Role change when synchronous mirroring is not operational

A more complex situation for role switching is when there is no communication between the two sites, either because of a network malfunction, or because the primary site is no longer operational.

The CLI command for this scenario is **mirror_change_role**. Because there is no communication between the two sites, the command should be issued on both sites concurrently, or at least before communication resumes. Otherwise, the sites will not be able to establish communication.

Switchover procedures differ depending on whether the Primary and Secondary volumes are connected or not. As a general rule:

- When the coupling is deactivated, it is acceptable to change the role on one side only, assuming that the other side will be changed as well before communication resumes.
- If the coupling is activated, but is either unsynchronized or non-operational due to a link error, an administrator must either wait for the coupling to be synchronized, or deactivate the coupling.
- On the Secondary volume, an administrator can change the role even if coupling is active. It is assumed that the coupling will be deactivated on the Primary volume and the role switch will be performed there as well in parallel. If not, a configuration error occurs on the original Primary volume.

Changing the Secondary volume to the Primary

The role of the Secondary volume can be changed to Primary, using Hyper-Scale Manager or CLI. After this switchover, the following takes effect:

- The Secondary volume is now the Primary volume.
- The coupling has the status of unsynchronized.
- The coupling remains in Standby mode, meaning that the remote mirroring is deactivated. This ensures an orderly activation when the role of the other site is switched.

The new Primary volume starts to accept write commands from local hosts. Because coupling is not active, in the same way as any Primary volume, it maintains a log of which write operations should be sent to the Secondary volume when communication resumes.

Typically, after switching the Secondary to the Primary volume, an administrator also switches the Primary to the Secondary volume, at least before communication resumes. If both volumes are left with the same role, a configuration error occurs.

Changing the Primary volume to the Secondary

When coupling is inactive, the primary machine can switch roles. After such a switch, the Primary volume becomes the Secondary.

Before switching roles, the Primary volume is inactive. Hence, it is in the unsynchronized state, and it might contain data that has not been replicated. Such data will be lost. When the Primary volume becomes Secondary, this data must be discarded to match the data on the peer volume, which is now the new Primary volume. In this case, an event is created, summarizing the size of the lost data.

Upon reestablishing the connection, the recovery volume (current Secondary, which was the Primary) will update the remote volume (new Primary) with this uncommitted data list to update, and it is the responsibility of the new Primary volume to synchronize these lists to the local volume (new Secondary).

I/O operations in synchronous mirroring

I/O operations are performed on the Primary and Secondary volumes across various configuration options.

I/O on the Primary volume

Read
: All data is read from the primary (local) site regardless of whether the system is synchronized.

Write
:
- If the coupling is operational, data is written to both the Primary and Secondary volumes.
- If the coupling is non-operational, data is written to the Primary volume only, and the Primary is aware that the Secondary is currently not synchronized.

I/O on the Secondary volume

The LUN of a Secondary volume can be mapped to remote hosts. In this case, the Secondary volume will be accessible to those remote hosts as *Read-only*.

These mappings are then used by remote hosts for Primary-Secondary role switchover. When the Secondary volume becomes the Primary, hosts can write to it on the remote site. When the Primary volume becomes a Secondary volume, it becomes *Read-only* and can be updated only by data replicated from the new Primary volume.

Read
: Data can be read from the Secondary volume like from any other volume.

Write
: In an attempt to write on the Secondary volume, the host will receive a volume read-only SCSI error.

Synchronization speed optimization

The storage system has two global parameters that limit the maximum rate used for initial synchronization and for synchronization after non-operational coupling.

These limits are used to prevent a situation where synchronization uses too much of the system or communication line resources, and hampers the host's I/O performance.

The values of these global parameters can be viewed by the user, but setting or changing them should be performed by an IBM® technical support representative.

Dynamic rate adaptation

The storage system provides a mechanism for handling insufficient bandwidth and external connections whenever remote mirroring is used.

The mirroring process replicates data from one site to the other. To accomplish this, the process depends on the availability of bandwidth between the local and remote storage systems. The mirroring synchronization rate parameter determines the bandwidth that is required for a successful mirroring.

You can request that an IBM technical support representative manually modify this parameter. To define its value, the IBM technical support representative should take into account the availability of bandwidth for the mirroring process, where the storage system adjusts itself to the available bandwidth.

The storage system prevents I/O timeouts through continuously measuring the I/O latency. Excessive incoming I/Os are queued until they can be submitted. The mirroring rate dynamically adapts to the number of queued incoming I/Os, allowing for a smooth operation of the mirroring process.

Implications on volume and snapshot management

When using sync mirroring, the default behavior of volumes and snapshots changes in order to protect the mirroring operation, as follows:

- Renaming a volume changes the name of the last-consistent and most updated snapshots.
- Deleting all snapshots does not delete the last-consistent and most updated snapshots.
- Resizing a Primary volume automatically resizes its Secondary volume.
- A Primary volume cannot be resized when the link is down.
- Resizing, deleting, and formatting are not permitted on a Secondary volume.
- A Primary volume cannot be formatted. If a Primary volume must be formatted, an administrator must first deactivate the mirroring, delete the mirroring, format both the Secondary and Primary volumes, and then define the mirroring again.
- Secondary or Primary volumes cannot be the target of a copy operation.
- Locking and unlocking are not permitted on a Secondary volume.
- The last-consistent and most updated snapshots cannot be unlocked.
- Deleting is not permitted on a Primary volume.
- Restoring from a snapshot is not permitted on a Primary volume.
- Restoring from a snapshot is not permitted on a Secondary volume.
- A snapshot cannot be created with the same name as the last-consistent or most updated snapshot.

Coupling synchronization process

When a failure condition has been resolved, remote mirroring begins the process of synchronizing the coupling. This process updates the Secondary volume with all the changes that occurred while the coupling was not operational.

Figure 1 shows the various coupling states, together with the actions that are performed in each state.

Figure 1. Coupling states and actions

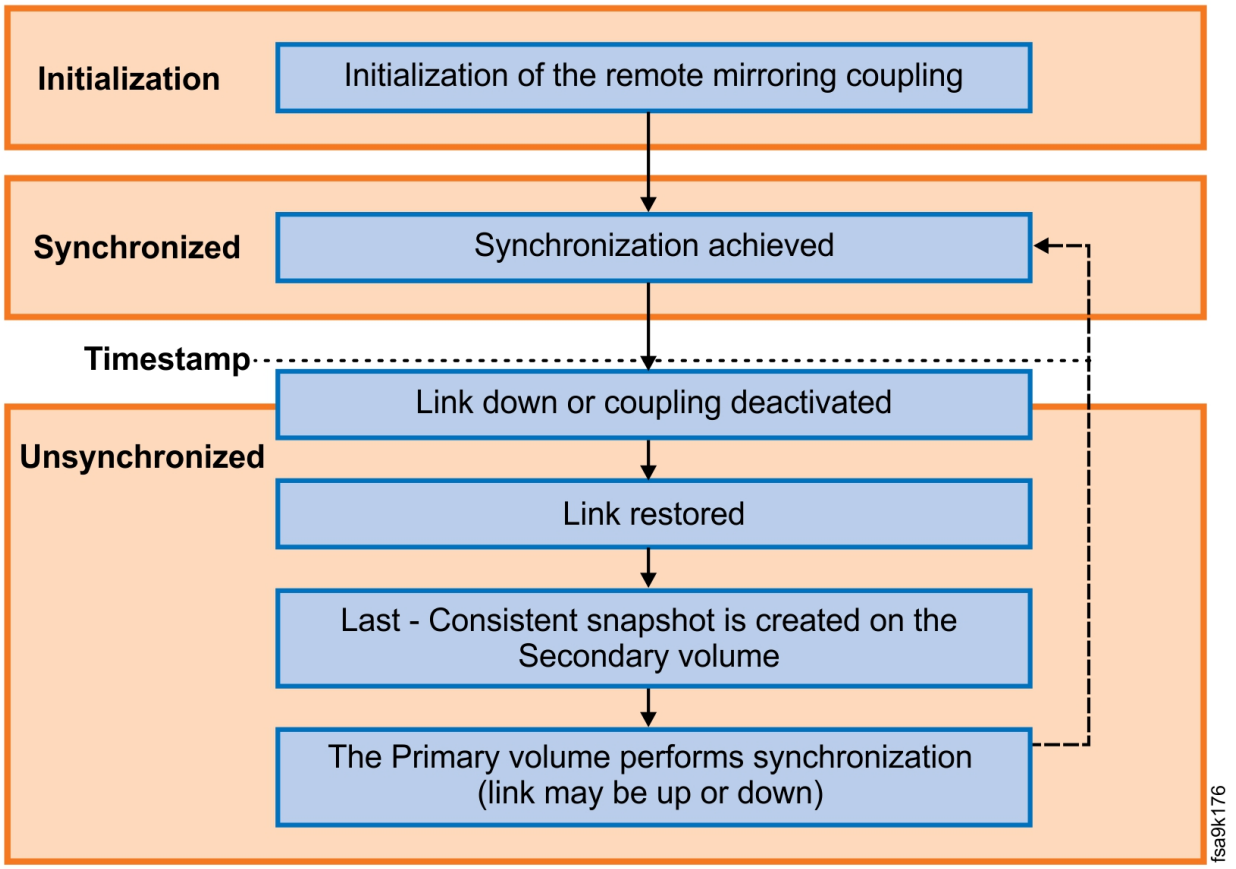

The following list describes each coupling state:

Initialization
: The Secondary volume has a *Synchronization* status of *Initialization*. During this state, data from the Primary volume is copied to the Secondary volume.

Synchronized
: This is the working state of the coupling, where the data in the Secondary volume is consistent with the data in the Primary volume.

Timestamp
: When a link is down, or when a coupling is deactivated, a timestamp needs to be taken. After the timestamp is taken, the state changes to *Timestamp*, and stays so until the link is restored, or the coupling is reactivated.

Unsynchronized
: Remote mirroring is recovering from a communications failure or deactivation. The Primary and Secondary volumes are being synchronized.

Coupling recovery

When remote mirroring recovers from a non-operational coupling, the following actions take place:

- If the Secondary volume is in the *Synchronized* state, a last-consistent snapshot of the Secondary volume is created and named with the string `secondary-volume-time-date-consistent-state`.
- The Primary volume updates the Secondary volume until it reaches the *Synchronized* state.
- When all couplings that mirror volumes between the same pair of systems are synchronized, the Primary volume deletes the special snapshot.

Uncommitted data

For best-effort coupling, when the coupling is in *Unsynchronized* state, the system must track which data in the Primary volume has been changed, so that these changes can be committed to the Secondary when the coupling becomes operational again.

The parts of the Primary volume that must be committed to the Secondary volume and must be marked are called *uncommitted data*.

Constraints and limitations

The following constraints and limitations apply to the synchronization process:

- The size, part, or time-to-synchronize are relevant only if the synchronization status is *Unsynchronized*.
- The last-secondary timestamp is only relevant if the coupling is *Unsynchronized*.

Synchronous mirroring of consistency groups

Mirroring can be applied to whole consistency groups.

The following restrictions apply:

- All volumes in a consistency group have the same role, either Primary, or Secondary
- All mirrors in a consistency group are between the same two systems

Asynchronous remote mirroring

Asynchronous mirroring enables high availability of critical data by asynchronously replicating data updates from a primary storage peer to a remote, secondary peer.

The relative merits of asynchronous and synchronous mirroring are best illustrated by examining them in the context of two critical objectives:

- Responsiveness of the storage system
- Currency of mirrored data

With synchronous mirroring, host writes are acknowledged by the storage system only after being recorded on both peers in a mirroring relation. This yields high currency of mirrored data (both mirroring peers have the same data), yet results in less than optimal system responsiveness because the local peer cannot acknowledge the host write until the remote peer acknowledges it. This type of process incurs latency that increases as the distance between peers increases, but both peers are synchronized (Figure 1).

Asynchronous mirroring (Figure 2) is advantageous in situations that warrant replication between distant sites because it eliminates the latency inherent to synchronous mirroring, and might lower implementation costs. Careful planning of asynchronous mirroring can minimize the currency gap between mirroring peers, and can help realize better data availability and cost savings.

Figure 1. Synchronous remote mirroring concept

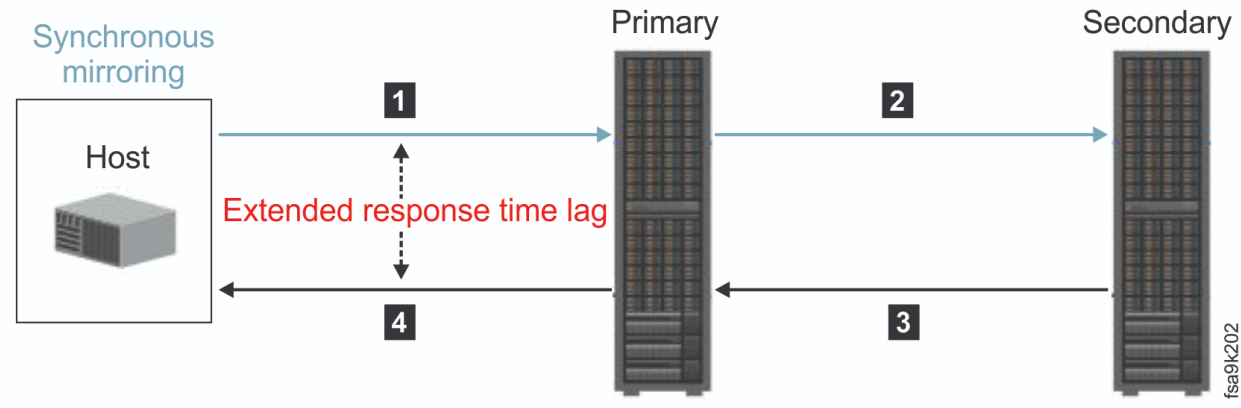

Figure 2. Asynchronous mirroring - no extended response time lag

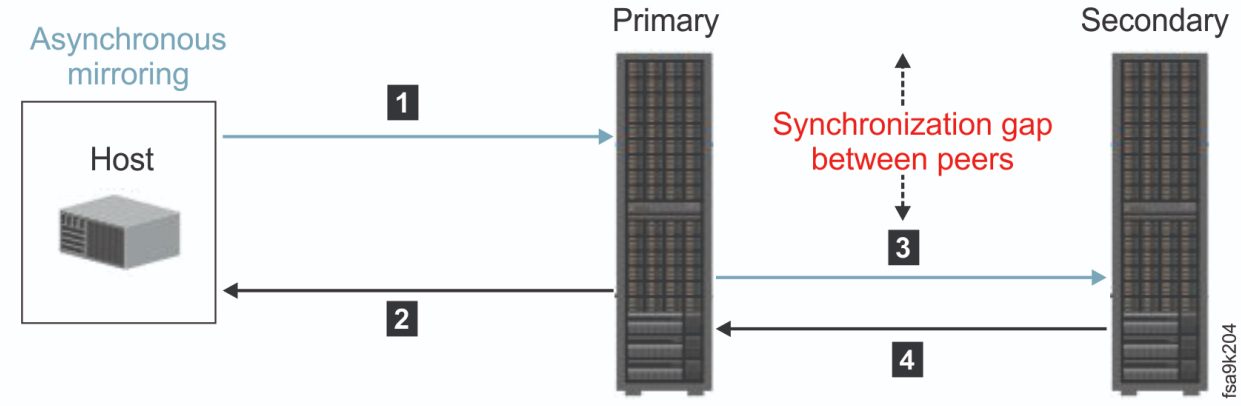

Note: Synchronous mirroring is covered in Synchronous remote mirroring.
For the information on the cross-system use of asynchronous remote mirroring, see the feature availability matrix.

- **Asynchronous mirroring highlights**
  The following are highlights of the asynchronous mirroring capability.
- **Snapshot-based technology in asynchronous mirroring**
  A snapshot-based technology for asynchronous mirroring facilitates concurrent mirrors with different recovery objectives.
- **Mirroring-special snapshots**
  The status and scope of the synchronization process is determined through the use of snapshots.
- **Disaster recovery scenarios in asynchronous mirroring**
  A disaster is a situation where one of the sites (either Primary or Secondary) fails, or the communication between the Primary site and the Secondary site is lost.

Asynchronous mirroring highlights

The following are highlights of the asynchronous mirroring capability.

Advanced snapshot-based technology
: Asynchronous mirroring is based on IBM snapshot technology, which streamlines implementation while minimizing impact on general system performance. The technology leverages functionality that supports mirroring of complete systems, translating to hundreds of mirrors. For a detailed description, see Snapshot-based technology in asynchronous mirroring.

Cross-generation asynchronous replication between A9000 systems and XIV Gen3
: Cross-generation asynchronous replication between your A9000 systems and XIV Gen3 systems is simplified by the IBM Hyper-Scale Manager user-interface, allowing you to leverage your XIV Gen3 investment, and lower the cost of data protection and disaster recovery.

Mirroring of consistency groups
: The storage system supports definition of mirrored consistency groups. In a consistency group, the mirrored image of the volumes at the remote site is kept consistent to the same point in time for all the volumes in the group. This is highly advantageous to enterprises, because it facilitates easy management of replication for all volumes that belong to a single consistency group. This also enables streamlined restoration of consistent volume groups from a remote site upon unavailability of the primary site.

Automatic and manual replication
: Asynchronous mirrors can be assigned a user-configurable schedule for automatic, interval-based replication of changes, or can be configured to replicate changes upon issuance of a manual (or scripted) user command. Automatic replication allows you to establish crash-consistent replicas, whereas manual replication allows you to establish application-consistent replicas, if required. You can combine both approaches, because you can define mirrors with a scheduled replication and issue manual replication jobs for these mirrors as needed.

Multiple RPOs (Recovery Point Objectives) and multiple schedules
: Asynchronous mirroring enables each mirror to be specified a different RPO, rather than forcing a single RPO for all mirrors. This can be used to prioritize replication of some mirrors over others, potentially making it easier to accommodate application RPO requirements, as well as bandwidth constraints.

Flexible and independent mirroring intervals
: Asynchronous mirroring supports schedules with intervals ranging between 20 seconds and 12 hours. Moreover, intervals are independent from the mirroring RPO. This enhances the ability to fine tune replication to accommodate bandwidth constraints and different RPOs.

Flexible pool management
: Asynchronous mirroring enables the mirroring of volumes and consistency groups that are stored in thin provisioned pools. This applies to both mirroring peers.

Bi-directional mirroring
: A storage system can host multiple mirror sources and targets concurrently, supporting over a thousand mirrors per system. Furthermore, any given storage system can have mirroring relations with several other storage systems. This enables enormous flexibility when setting mirroring configurations.
The number of systems with which the storage system can have mirroring relations is specified outside, in the Data Sheet.

Concurrent synchronous and asynchronous mirroring
: The storage system can concurrently run synchronous and asynchronous mirrors.

Easy transition between peer roles
: Mirror peers can be easily changed between Primary and Secondary.

Easy transition from independent volume mirrors into consistency group mirror
: The asynchronous mirroring allows for easy configuration of consistency group mirrors, easy addition of mirrored volumes into a mirrored consistency group, and easy removal of a volume from a mirrored consistency group while preserving mirroring for such volume.

Control over synchronization rates per target
: The asynchronous mirroring implementation enables administrators to configure different system mirroring rates with each target system.

Comprehensive monitoring and events
: Storage systems generate events and monitor critical asynchronous mirroring-related processes to produce important data that can be used to assess the mirroring performance.

Easy automation via scripts
: All asynchronous mirroring commands can be automated through scripts.

Snapshot-based technology in asynchronous mirroring

A snapshot-based technology for asynchronous mirroring facilitates concurrent mirrors with different recovery objectives.

With asynchronous mirroring, write order on the Primary volume is not preserved on the Secondary volume. As a result, a snapshot taken of the Secondary volume at any moment is most likely inconsistent and therefore not valid. To ensure high availability of data in the event of a failure or unavailability of the Primary volume, it is imperative to maintain a consistent replica of the Primary volume that can ensure service continuity.

This is achieved through storage system snapshots. Asynchronous mirroring employs them to record the state of the Primary volume, and calculates the difference between successive snapshots to determine the data that needs be copied from the Primary to the Secondary volume as part of a corresponding replication process. Upon completion of the replication process, a snapshot is taken of the Secondary volume and reflects a valid replica of the Primary volume.

Below are select technological properties that explain how the snapshot-based technology helps realize effective asynchronous mirroring:

- The storage system supports a practically unlimited number of snapshots, which facilitates mirroring of complete systems with practically no limitation on the number of mirrored volumes supported
- Asynchronous mirroring implements memory optimization techniques that further maximize the performance attainable by minimizing disk access.

Mirroring-special snapshots

The status and scope of the synchronization process is determined through the use of snapshots.

The following two special snapshots are used:

most recent snapshot (MRS)
: This snapshot is the most recent snapshot of the Primary volume or consistency group, prior to the creation of a new replication process. This snapshot exists only on the Primary system.

last replicated snapshot (LRS)
: This is the most recent snapshot that is confirmed to have been fully replicated to the Secondary system. Both the Primary and the Secondary systems have this snapshot. On the Secondary system, the snapshot is taken upon completion of a replication process, and replaces any previous snapshot with that name. On the Primary system, the most recent snapshot is renamed to *last_ replicated* after the Secondary machine is confirmed to have a corresponding replica of the Primary system's most recent snapshot.

Figure 1. Location of special snapshots

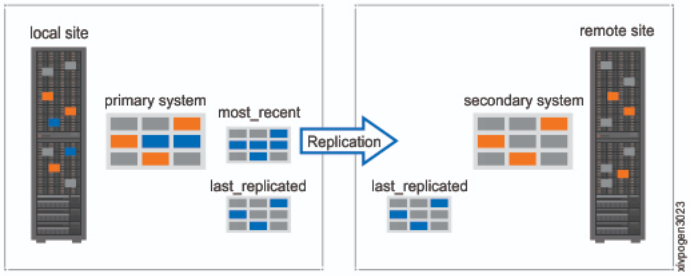

IBM FlashSystem A9000 or A9000R maintains three snapshots per mirror coupling: two on the Primary system and one on the Secondary. A valid (recoverable) state of the Primary volume is captured in the last replicated snapshot, and on an identical snapshot on the Secondary system. The most recent snapshot represents a recent state of the Primary volume that needs be replicated next to the Secondary system. The system determines the data to replicate by comparing the Primary volume's snapshots.

Disaster recovery scenarios in asynchronous mirroring

A disaster is a situation where one of the sites (either Primary or Secondary) fails, or the communication between the Primary site and the Secondary site is lost.

Asynchronous mirroring attains synchronization between the Primary and Secondary peers through a recurring data replication process called a *Sync Job*. Running at user-configurable schedules, the Sync Job takes the most recent snapshot of the Primary volume and compares this snapshot with the last replicated snapshot on the Secondary volume. The Sync Job then synchronizes the Primary volume data corresponding to these differences with the Secondary volume. At the completion of a sync job, a new last replicated snapshot is created on both the Secondary and Primary volumes.

Disaster recovery scenarios handle cases in which one of the snapshots mentioned above becomes unavailable. These cases are:

Unplanned service disruption

1 Failover
: Unplanned service disruption starts with a failover to the Secondary volume.
The Secondary is promoted and becomes the new Primary, serving host requests

2 Recovery
: Next, whenever the Primary volume and the link are restored, the replication is set from the promoted Secondary (the new Primary) onto the demoted Primary (the new Secondary).

Alternatively: No recovery
: If recovery is not possible, a new mirroring is established on the Secondary volume. The original mirroring is deleted and a new mirroring relation is defined.

3 Failback
: Following the recovery, the original mirroring configuration is reestablished. The Primary volume maintains its role and replicates to the Secondary volume.

Planned service disruption

1 Planned role switch
: Planned service disruption starts with a coordinated demotion of the Primary to the Secondary, while the Secondary is promoted to become the new Primary. The promoted Secondary volume serves host requests, and replicates to the demoted Primary. On the host side, the host is disconnected from the demoted Primary and connected to the new Primary.

2 Recovery
: Next, whenever the Primary volume and the link are restored, the replication is set from the promoted Secondary volume (the new Primary) onto the demoted Primary (the new Secondary volume).

2 Failback
: Following the recovery, the original mirroring configuration is reestablished. The Primary volume maintains its role and replicates to the Secondary volume.

Testing
: There are two ways to test the Secondary volume replica:

- Create a snapshot of an LRS snapshot on the Secondary volume. Then map a host to it and verify the data.
- Disconnect the host from the Primary volume, switch roles, and connect the host to the Secondary volume. This is a more realistic, but also a more disruptive test.

Note: Please contact IBM® Support in case of disaster or for any testing of disaster recovery, in order to get clear guidelines and to secure a successful test.

High availability with HyperSwap

HyperSwap® delivers highly-available, non-disruptive storage service, through partial or complete system failures and disasters, in the same data center and between metro-distant data centers.

HyperSwap high availability is based on active-active pairing of storage systems per volume or per consistency group. Each volume or consistency group pair uses synchronous replication to keep both systems updated at all times.

When certain conditions apply, an automatic and completely transparent failover is performed so that the applications experience no downtime. As soon as the actual failure is recovered, the pair is automatically resynchronized.

As in other high availability solutions, HyperSwap requires a quorum witness component to avoid split-brain situations. HyperSwap Quorum Witness is constantly monitoring the status of the related storage systems, and, if necessary, acts as a tiebreaker for conflict resolution.

The HyperSwap solution relies on Asymmetrical Logical Unit Access (ALUA) support to inform the host about the optimized paths to the storage system, and minimize I/O latency.

FlashSystem A9000 and FlashSystem A9000R HyperSwap capability does not require additional special hardware or software, and does not require any additional licensing.

The system supports a maximum of 1536 mirroring relations.
Important: FlashSystem A9000R systems support a maximum of 3072 mirror couplings. These couplings are shared between all and any combinations of synchronous, asynchronous, HyperSwap and Multi-site relations. The storage system high availability feature allows for a maximum of 1536 HyperSwap pairs. These pairs are a part of the total allowance of 3072 mirror couplings that the system supports.
For the information on the cross-system use of HyperSwap, see the feature availability matrix.

- **Design guidelines**
  To properly configure and implement high availability with IBM HyperSwap, it is important to understand its principles and key components.
- **Operating a HyperSwap-based High Availability solution**
  The following topics describe how to operate a HyperSwap-based High Availability solution.

Design guidelines

To properly configure and implement high availability with IBM HyperSwap, it is important to understand its principles and key components.

This section provides general information about the IBM HyperSwap main components, configuration topology, and supported networking. It also describes the HyperSwap volume, and explains in which situations an automatic failover is expected.

- **Configuration**
  This sub-section provides general information about the IBM® HyperSwap solution components, the concept of a failure domain, the HyperSwap volume, the Quorum Witness, and configuration requirements.
- **Topologies**
  This section surveys various configurations of hosts, storage systems and Quorum Witnesses for HyperSwap volumes.
- **Automatic failover scenarios**
  IBM FlashSystem® A9000 and IBM FlashSystem A9000R HyperSwap solution is designed to maintain non-disruptive host data access.

Configuration

This sub-section provides general information about the IBM® HyperSwap® solution components, the concept of a failure domain, the HyperSwap volume, the Quorum Witness, and configuration requirements.

Important: Considerations must be made when connecting 8 or more FlashSystem A9000 and/or A9000R systems to a HyperSwap Quorum Witness. For more information, contact IBM Support.
A minimal HyperSwap solution consists of:

- Two IBM FlashSystem® A9000 or A9000R storage systems, interconnected for synchronous replication via Fibre Channel
- HyperSwap-protected hosts, each connected to both systems via iSCSI or Fibre Channel
- Quorum Witness software, installed on a VM or a physical host, with TCP-IP connectivity to both systems.

The paired systems maintain one or more HyperSwap relations between them. Each relation facilitates one HyperSwap volume (see HyperSwap volumes and consistency groups).

Larger HyperSwap solutions are possible, since every IBM FlashSystem A9000 and A9000R system can have HyperSwap relations with multiple other systems.

- **HyperSwap volumes and consistency groups**
  A HyperSwap volume is implemented as a pair of two volumes with identical SCSI attributes, one on each system.
- **Quorum Witness**
  Every clustering solution requires a Quorum Witness component that can be consulted by the cluster members at any time to avoid split-brain situations.
- **Failure domain**
  A failure domain encompasses all the elements potentially affected by a single failure.

HyperSwap volumes and consistency groups

A HyperSwap® volume is implemented as a pair of two volumes with identical SCSI attributes, one on each system.

These volumes should be always kept synchronized. From the host perspective, the two volumes are a single volume, and I/O – both reads and writes - can be served from either system, depending on the path used by the host. In other words, a HyperSwap solution allows a host to have active-active access to the same data on two systems. Figure 1 shows the replication flow between the Primary and the Secondary volumes.

Figure 1. HyperSwap volume replication

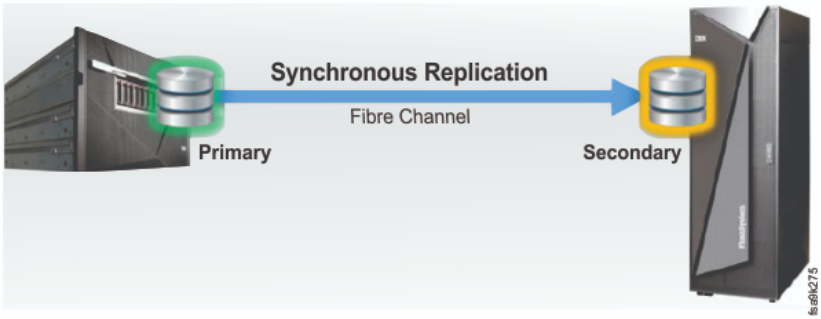

- To be identical, a pair of volumes that constitute a HyperSwap volume have identical SCSI identity and I/O related attributes - size, locks and reservations. Each storage system maps the paired volumes to the host separately, but the host perceives them as a single volume. This makes transitions, such as automatic failover and manual failback, transparent both to the hosts and the applications running on them.
- To be synchronized, the peer-systems are interconnected for synchronous replication, and a HyperSwap relation is established between the peer-volumes. One of these volumes is initially designated as Primary, and the other is designated as Secondary. As opposed to mirroring, the replication between the Primary and Secondary volumes is bi-directional, to allow read and write I/O to be served on either volume. The purpose of the Primary/Secondary designation is to optimize latency: the Primary volume should be co-located with the hosts that generate most of the I/O.
- The actual roles – Primary or Secondary - performed by the volumes at any given moment may differ from their original role designations, as a result of a manual role change or an automatic failover.

Multiple HyperSwap volumes can exist between any HyperSwap-paired systems. Moreover, HyperSwap volumes can be grouped into FlashSystem A9000/A9000R consistency groups. In such a case, HyperSwap actions will be applied to the entire consistency group, that is, to all the HyperSwap volumes in it.

The following restrictions apply to HyperSwap consistency groups:

- All the volumes in a HyperSwap consistency group must be HyperSwap volumes
- All the volumes in a HyperSwap consistency group must have the same target.

Important: When using IBM® Hyper-Scale Mobility on Storage Area Network (SAN) Boot Volumes, the Proxy targets need to be replaced with the Owner targets on the adapter BIOS. In some cases, this can be done online from the operating system, by using vendor's online tools for x86 servers (for example, Qlogic QConvergeConsole), or by using specific commands to edit the boot paths order for Unix servers. In other cases, manual changes of the adapter BIOS are required.

Quorum Witness

Every clustering solution requires a Quorum Witness component that can be consulted by the cluster members at any time to avoid split-brain situations.

Important: Considerations must be made when connecting 8 or more FlashSystem A9000 and/or A9000R systems to a HyperSwap Quorum Witness. For more information, contact IBM Support.
In the HyperSwap® solution, this function is performed by the IBM Spectrum Accelerate Family HyperSwap Quorum Witness software application. It allows FlashSystem A9000 & A9000R arrays to determine in any conditions and at any time, for each HyperSwap volume, which single array should exclusively own the Primary volume instance.
Connectivity between a Quorum Witness and storage systems is established via TCP-IP. For more information, see the Topologies section.

When the Quorum Witness is down for any reason, there is no impact on HyperSwap active-active data access, and various failure scenarios will still be accommodated without disruption. However, during that period, automatic failover cannot be applied, since the FlashSystem A9000 and A9000R systems have no other way to ensure that a failover will not result in a split-brain situation. Therefore when the Quorum Witness is down, the risk of downtime is elevated.

To minimize Quorum Witness downtime, it should be continuously monitored, and any issues should be addressed immediately. To further minimize the risk of Quorum Witness downtime, it can be deployed as a highly-available VM on a VMware vSphere cluster, using VMware High Availability.
Important: Quorum Witness cannot reside on the same set of FlashSystem A9000 and/or A9000R systems of which it serves as part of a HyperSwap configuration.

Failure domain

A failure domain encompasses all the elements potentially affected by a single failure.

For example, an earthquake or power-grid failure can potentially take down a whole datacenter. The datacenter is therefore a failure domain, with regards to earthquake and power-grid failures. To protect against such failures, the infrastructure can be divided between two geographically-separate data centers that are not using the same power grid. Each data center is considered a failure domain, and if one of them fails, the other can replace it. Failure domains can also be

defined within a single datacenter, depending on the kind of failure they protect against. For example, to protect against overheating, the datacenter may have multiple cooling systems. The area that is protected by each cooling system will be a failure domain, from cooling perspective.

Therefore, for every two systems that are HyperSwap®-paired, the best practice is to separate and spread the two systems and the Quorum Witness across three failure-domains to prevent the situation where more than one of these elements is down. The highest availability level is obtained when these three failure domains are 3 geographically separated sites, with separate power and network resources.
Important: Quorum Witness cannot reside on the same set of FlashSystem A9000 and/or A9000R systems of which it serves as part of a HyperSwap configuration.

Topologies

This section surveys various configurations of hosts, storage systems and Quorum Witnesses for HyperSwap® volumes.

- **Volume level perspective**
  This section focuses on a typical configuration of hosts, storage systems and a Quorum Witness for one HyperSwap volume.
- **System level perspective**
  Since every system can contain a mix of multiple Primary and Secondary volumes and have HyperSwap relations with multiple systems, it is interesting to consider topologies at system level, with multiple HyperSwap volumes.
- **Other topologies**
  Since every system can have HyperSwap relations with up to ten other systems, and serve a mix of Primary and Secondary volumes, other topologies are possible as well.

Volume level perspective

This section focuses on a typical configuration of hosts, storage systems and a Quorum Witness for one HyperSwap® volume.

Typically, the host and the storage system initially designated for Primary volume are located at the same site, and the Quorum Witness is deployed at a separate third site.

In Figure 1, the choice of FlashSystem A9000R on one side, and FlashSystem A9000 on the other side is not necessarily typical, it is only used to emphasize that it is a valid configuration.

Figure 1. Typical HyperSwap high availability configuration

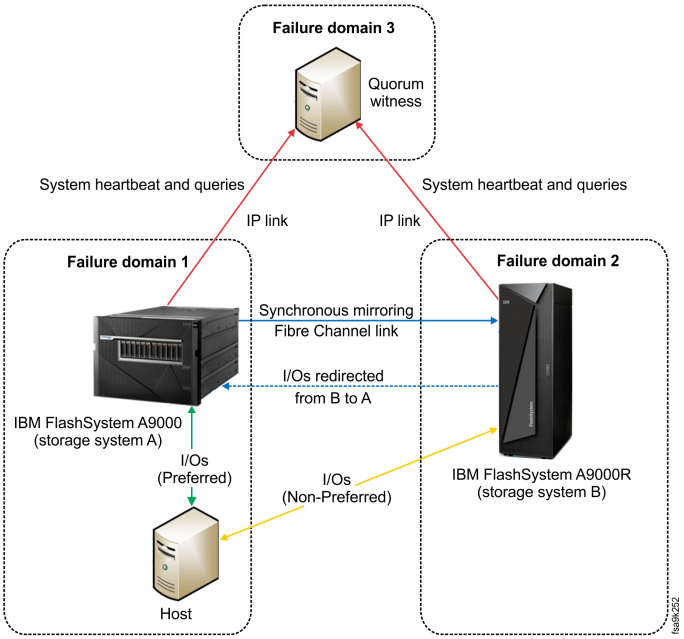

Host-to-storage-system paths, hereafter referred to as port groups, are optimized using Asymmetric Logical Unit Access (ALUA) support from the multipath driver. By assigning proper ALUA states (Preferred or Non-Preferred), the storage system directs the host multipath driver which paths are preferred, to minimize I/O latency:

- Port groups to the system that currently owns the Primary volume are automatically marked as Active/Preferred
- Port groups to the system that currently owns the Secondary volume are automatically marked as Active/Non-Preferred.

As a result, Active/Preferred port groups receive the bulk of the I/O traffic. Any remaining I/O traffic is directed to the Active/Non-Preferred port groups, and writes are then forwarded to the Primary volume.

When the HyperSwap volume is activated, the Secondary volume is not synchronized, and read requests are redirected to the Primary volume until synchronization has completed. When the volumes are synchronized, the system that owns the Secondary volume serves read requests locally.

If the system that owns the Secondary volume is unable to perform I/O, the Secondary volume port group state changes to Unavailable. Such a change would typically be a result of a connectivity failure. The Primary volume remains active, therefore no automatic failover is needed. As soon as connectivity is restored, the

volumes will be re-synchronized automatically.

If the system that owns the Primary volume is unable to perform I/O, the Primary volume port group state changes to Unavailable. As soon as the system that owns the Secondary volume receives the corresponding notification from the Quorum Witness, it performs a transparent failover, by which the Secondary volume assumes the Primary role. When the system that was originally designated for the Primary volume is restored, recovery must be performed manually via CLI or via IBM® Hyper-Scale Manager, which involves switching the roles of the peer volumes and re-activating the HyperSwap relation between them.

System level perspective

Since every system can contain a mix of multiple Primary and Secondary volumes and have HyperSwap® relations with multiple systems, it is interesting to consider topologies at system level, with multiple HyperSwap volumes.

Here are a few examples, where green denotes Primary volumes, and yellow denotes Secondary volumes.

Conventional HyperSwap topology

In a conventional Disaster Recovery topology (that is, where one system is located in a site that is designated by the storage operations for disaster recovery) one storage system owns the volumes designated as Primary, and the other system owns the volumes designated as Secondary, as shown in Figure 1.

Figure 1. Conventional HyperSwap topology

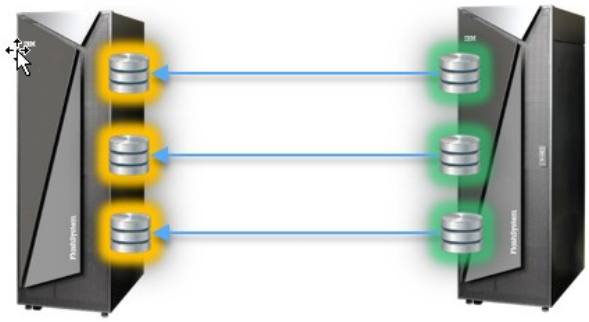

Star-shaped HyperSwap Disaster Recovery topology

In a star-shaped HyperSwap Disaster Recovery topology, where one of the sites is designated by the storage operations as a Disaster Recovery site, a single storage system in the Disaster Recovery site is dedicated to simultaneously serving multiple other systems, as shown in Figure 2.

Figure 2. Star-shaped HyperSwap Disaster Recovery topology

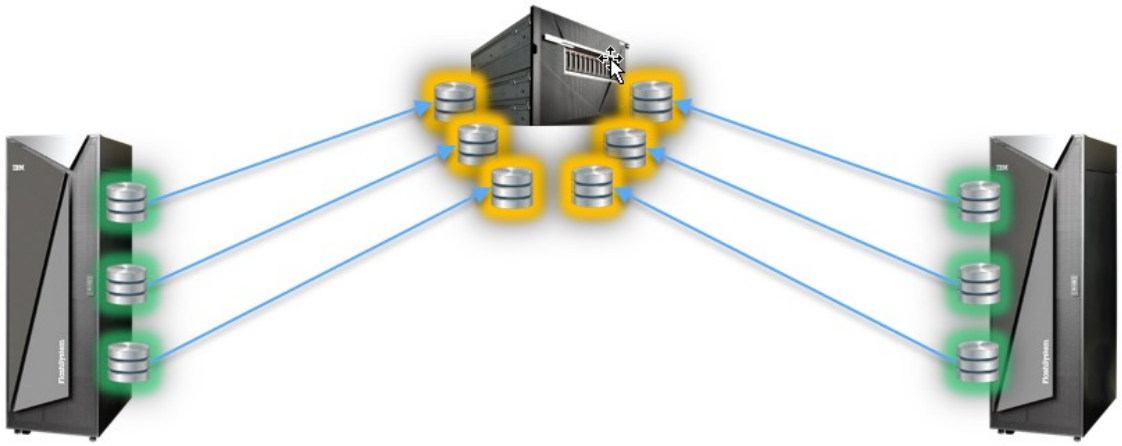

Symmetrical HyperSwap topology

In a symmetrical HyperSwap topology, both systems have volumes designated as Primary and Secondary, depending on the preferred location of the application, as shown in Figure 3.

Figure 3. Symmetrical HyperSwap topology

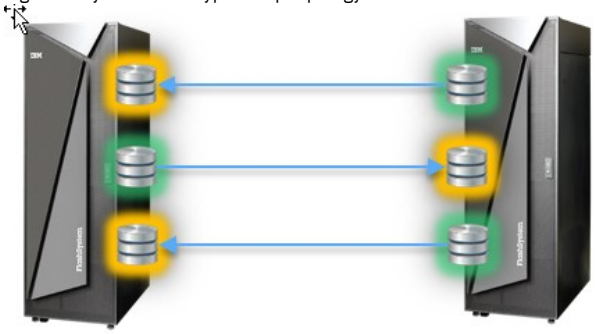

Other topologies

Since every system can have HyperSwap® relations with up to ten other systems, and serve a mix of Primary and Secondary volumes, other topologies are possible as well.

Two types of configuration of HyperSwap topology are possible based on how hosts are connected to storage systems.

In a uniform configuration, each host can access both the Primary and Secondary volumes. The uniform configuration is the best practice to protect a host from data access problems:

Figure 1. HyperSwap configuration types: uniform host connectivity

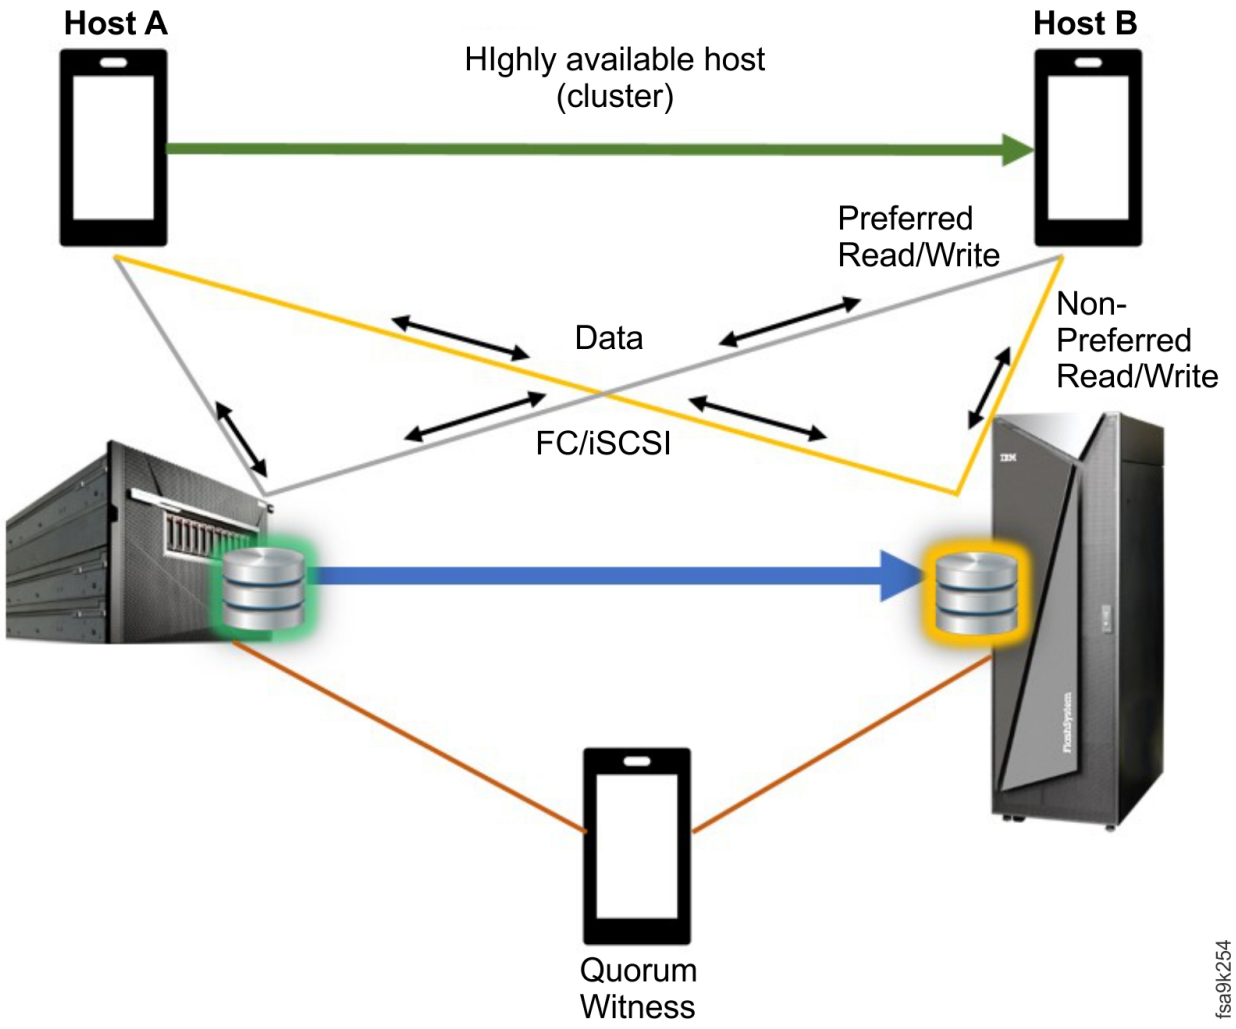

In a non-uniform configuration, the storage high availability relies on the server detection and failover of the application to the server with access to active storage. A non-uniform configuration can be used when the host is part of a cluster, can fail over to another host in the cluster, and the other host is connected to the peer system. This configuration is less costly from network perspective. However, it relies on a complex failover orchestration between host-clusters and the FlashSystem A9000 or FlashSystem A9000R systems, which is not needed in a uniform configuration.

Figure 2. HyperSwap configuration types: non-uniform host connectivity

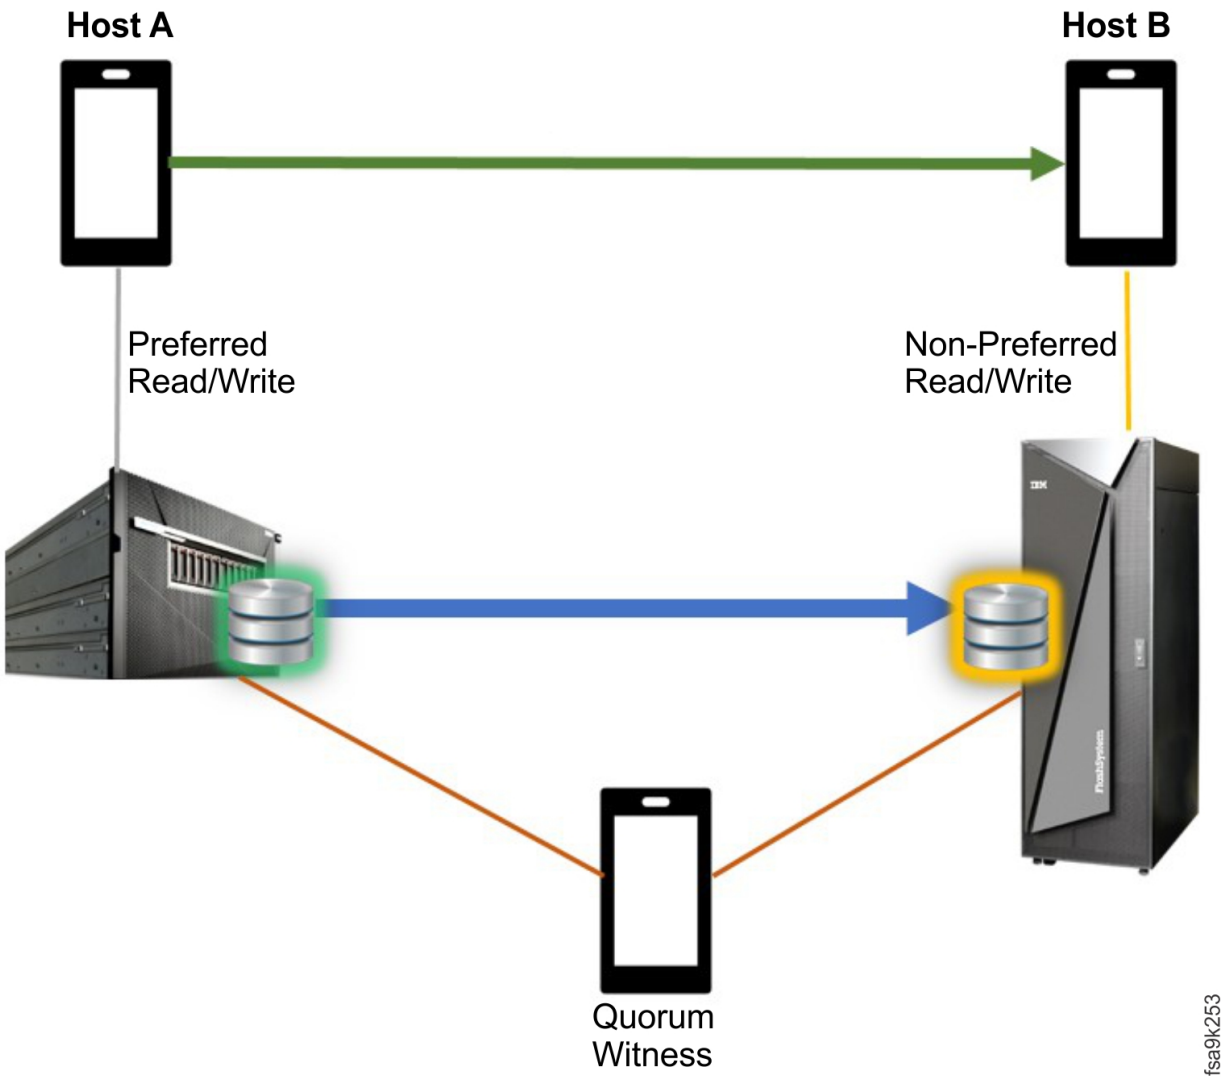

Automatic failover scenarios

IBM FlashSystem® A9000 and IBM FlashSystem A9000R HyperSwap® solution is designed to maintain non-disruptive host data access.

The HyperSwap decision-making process guarantees non-disruptive data access, while avoiding unnecessary failover.

To achieve this, the HyperSwap logic considers two aspects:

- System and Quorum Witness failure identification
- Assignment of Primary and Secondary roles to HyperSwap volumes.

The Quorum Witness allows the peer systems to avoid making erroneous or even conflicting decisions regarding the role of each volume in a HyperSwap volume pair. For instance, if connectivity between them is disrupted, the system that owns the Secondary volume may assume that its peer system has failed, and attempt to initiate an automatic failover in order to become the owner of the Primary volume. If this attempt succeeds while the system that owns the Primary volume is still serving I/O, both systems will claim that they own the Primary volume. The Quorum Witness protects against such split-brain scenarios. If the Primary volume owner system is still communicating with the Quorum Witness, the peer system will be aware of it. Instead of attempting a failover, it will render its own port groups *Unavailable*. Replication will stop and resume only after the connectivity problem is resolved.

Operating a HyperSwap-based High Availability solution

The following topics describe how to operate a HyperSwap-based High Availability solution.

- **Establishing a HyperSwap relation**
  A HyperSwap relation is similar in many ways to a mirroring relation, but it has extra attributes and states, that are related to the HyperSwap volume and high availability logic.
- **Monitoring a HyperSwap volume**
  The IBM Hyper-Scale Manager can display the current status of a HyperSwap volume at any time.

- **Modifying a HyperSwap volume**
  Any change that affects the SCSI attributes of the HyperSwap volume is automatically carried out on both the Primary and Secondary volumes.
- **High Availability snapshots**
  HyperSwap supports the creation of identical single-point-in-time snapshots at the local and remote systems.
- **Failure detection and recovery scenarios**
  The IBM HyperSwap feature is designed to detect a failure in a timely manner and automatically choose the most appropriate course of action for fast and comprehensive recovery.

Establishing a HyperSwap relation

A HyperSwap relation is similar in many ways to a mirroring relation, but it has extra attributes and states, that are related to the HyperSwap volume and high availability logic.

Configuring a HyperSwap environment

Before provisioning the first HyperSwap volume, a HyperSwap environment must be configured.

Configuration of a HyperSwap environment begins with a Quorum Witness deployment, as described in the *Quorum Witness User Guide*. In IBM Hyper-Scale Manager, define the same Quorum Witness on all the participating peer systems using one simple configuration, as seen inFigure 1.

Figure 1. Quorum Witness information through the Hyper-Scale Manager

When connectivity with the Quorum Witness is established, the system registers itself with the Quorum Witness, and the Quorum Witness is considered active by that system. Each peer system continuously verifies that it has connectivity with the target system and the Quorum Witness, and that the target system has connectivity with the same Quorum Witness.

Creating a new HyperSwap relation

As a prerequisite for creating a HyperSwap relation, the peer systems verify that:

- They are both registered on the same Quorum Witness
- The remote system is identified via the Quorum Witness as *healthy*
- The overall connectivity state is suitable for High Availability

If the checks are completed successfully, define the Primary and Secondary volumes and Consistency Groups to create a new HyperSwap relation. As explained in Configuration, the Secondary volume in a HyperSwap relation adopts the identity of the Primary volume.

Creating a HyperSwap relation from an existing synchronous mirror

A synchronous mirror relation can be transformed into a HyperSwap relation, if the following requirements are met:

- All the required configuration has completed successfully: the Quorum Witness is configured and activated, and the remote target configuration is updated
- The Secondary volume is unmapped
- The Quorum Witness is active and connected
- Both mirror volume names are identical.
  If a Consistency Group is converted into a HyperSwap Consistency Group, the names of all the volumes in the Secondary Consistency Group must be identical with the names of their peers in the Primary Consistency Group.

The transformation can be carried out only when the mirror is active. It does not disrupt the mirroring, and the data on the Secondary volume remains intact.

Activating a HyperSwap relation

When a HyperSwap relation is activated, the Quorum Witness configurations are verified. After the verification has completed successfully, the Primary volume starts synchronizing data to the Secondary volume. When the relation is activated, the Secondary volume becomes available, and active-active access is enabled.

The standard mirroring initialization methods are applicable to HyperSwap relations:

- Online: This method is the default one. It uses an inter-site link to replicate the Primary volume's initial state to the Secondary volume, starting once the mirror is first activated.

- Offline: Initialization of the Secondary volume is not done by replicating the Primary volume's initial image, but rather by Offline Initialization. In other words, it restores to the Secondary volume a mirror image that was backed up from the Primary volume. Once the relation is activated, the contents of the volumes are compared, and only modified data is synchronized over the wire. This process is usually much faster than online initialization, especially if the link throughput is limited.

For more information about mirroring initialization, see an IBM Redpaper publication IBM XIV Storage System Multi-Site Mirroring.
During initialization, the HyperSwap relation status is *Initializing*.

HyperSwap relation between Consistency Groups

You can establish a HyperSwap relation between HyperSwap Consistency Groups.

A HyperSwap Consistency Group is a Consistency Group that has a HyperSwap relation with a remote Consistency Group. To be eligible for the addition to a HyperSwap Consistency Group, the volume must conform to the following requirements:

- The volume is a HyperSwap volume
- The volume role (Primary or Secondary) matches the role of the Consistency Group
- The volume is connected to the same target system or peer system as the Consistency Group
- The HyperSwap relation is synchronized

Any later change to the Consistency Group, whether automatic or manual, affects all the volumes it contains. In particular, an automatic failover is carried out only if all the volumes in the Consistency Group are ready for the automatic failover.

Monitoring a HyperSwap volume

The IBM Hyper-Scale Manager can display the current status of a HyperSwap volume at any time.

The HyperSwap volume status is displayed in the Volume Availability window. The status depends on:

- Peer systems connectivity
- Availability of the Quorum Witness
- Ability of both Primary and Secondary volumes to handle I/O (*active-active* configuration)
- Synchronization between the Primary and Secondary volumes.

In the figure below, all the prerequisites for automatic failover readiness are fulfilled, as indicated by green check-boxes. Consequently, the Auto. Failover Capability box is green, which means that this volume is ready to perform an automatic failover when necessary.

Figure 1. HyperSwap volume status through Hyper-Scale Manager

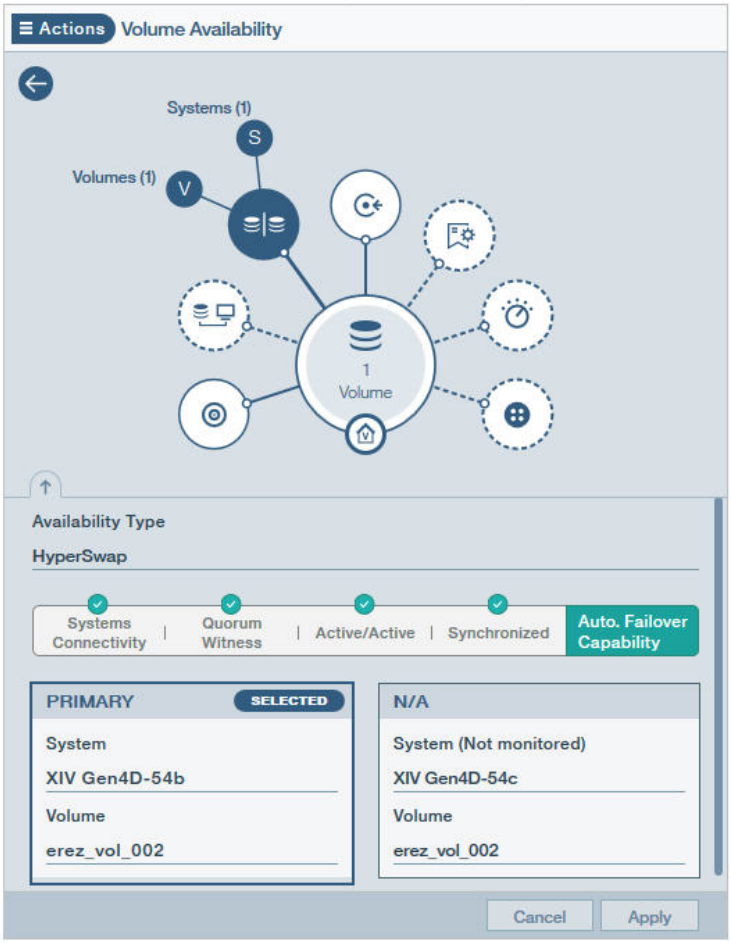

In addition, you can view the HyperSwap volume status in the Group by Availability column of the Volumes table.

Pointing at a volume marked as H in the Group by Availability column, displays a pop-up window with the volume's HyperSwap availability status, such as shown in Figure 2.

Figure 2. HyperSwap volume status in the Volumes table through Hyper-Scale Manager

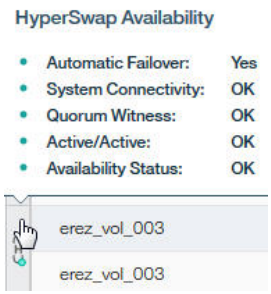

Modifying a HyperSwap volume

Any change that affects the SCSI attributes of the HyperSwap volume is automatically carried out on both the Primary and Secondary volumes.

Volume attributes include WWN, name, size, locking status, and SCSI reservation. Therefore, having a healthy connectivity is mandatory for many HyperSwap volume related actions.

High Availability snapshots

HyperSwap supports the creation of identical single-point-in-time snapshots at the local and remote systems.

A snapshot is a logical volume reflecting the contents of a given source volume at a specific point-in-time. FlashSystem A9000 and A9000R storage systems use snapshots to identify inconsistencies that may arise between data updates. Snapshot taking and management are described in detail in Volume function and lifecycle.

In regular mirroring, snapshots are created on the local system, as described in Synchronous mirroring operations.

In addition to regular snapshots that can be applied to the Primary or Secondary volumes, HyperSwap supports the creation of identical single-point-in-time snapshots on the local and remote systems. The operation is performed on the Primary volume, but the resulting snapshots are created on both the local and remote systems. The content of these snapshots is identical, but unlike the Primary and Secondary volumes, that have identical SCSI attributes, the SCSI identities of those co-created High Availability snapshots are different. Therefore the snapshots are independent of each other, and any snapshot operations - such as snapshot deletion or unlocking - are performed on the individual snapshots. Moreover, if one of the snapshots is unlocked, and writes are applied to it, the content of the snapshots stops being synchronized and is no longer identical.

Failure detection and recovery scenarios

The IBM HyperSwap feature is designed to detect a failure in a timely manner and automatically choose the most appropriate course of action for fast and comprehensive recovery.

In the most sensitive scenario, upon detecting that the system that owns the Primary volume (going forward referred to as System A) has failed, the system that owns the Secondary volume (going forward referred to as System B) should trigger an automatic failover.

The specific failover scenarios and actions are described in the following sections.
Important: Quorum Witness cannot reside on the same set of FlashSystem A9000 and/or A9000R systems of which it serves as part of a HyperSwap configuration.

- **System A - System B connectivity failure**
  If synchronous replication between Systems A and B is disrupted, System A remains solely responsible for serving I/O.
- **System A - Quorum Witness connectivity failure**
  If connectivity between System A and Quorum Witness is disrupted, System A will periodically attempt to recover connectivity. While Quorum Witness connectivity is not available, System B remains ready for automatic failover. Automatic failover will be triggered, if connectivity between the systems A and B fails as well.
- **System A - System B and System A - Quorum Witness connectivity failure or System A failure**
  If connectivity between System A and the Quorum Witness is disrupted, and connectivity between System A and System B is down, System A stops serving I/O and assumes that System B performed an automatic failover.
- **System B - Quorum Witness connectivity failure**
  If connectivity between System B and Quorum Witness is disrupted, System A will continue performing as usual. System B will periodically attempt to recover connectivity with the Quorum Witness, but readiness for automatic failover will be disabled.
- **System B - System A and System B - Quorum Witness connectivity failure or System B failure**
  If connectivity between System B and System A is disrupted, the replication is disrupted as well, and the Secondary volume goes out of sync. System A will then assume sole responsibility for serving I/O. Readiness for automatic failover will be disabled.
- **System A - Quorum Witness and System B - Quorum Witness connectivity failure**
  This scenario assumes that neither System A nor System B is connected to the Quorum Witness, while the System A - System B connectivity is in proper

order. Automatic failover readiness will be disabled. System B, however, continues serving I/O, and data on both systems remains synchronized.

- **System A I/O serving failure**
  If System A is functional, but is unable to serve I/O for any reason, it notifies System B of the failure, locks its HyperSwap volume, and sets the state of its port groups to *Unavailable*, thus blocking even I/O in flight. System B then automatically takes over the Primary volume and starts serving I/O.
- **Total connectivity failure**
  This scenario is very unlikely and practically impossible, if the storage arrays and Quorum Witness are located in separate sites.

System A - System B connectivity failure

If synchronous replication between Systems A and B is disrupted, System A remains solely responsible for serving I/O.

Figure 1. System A - System B connectivity failure

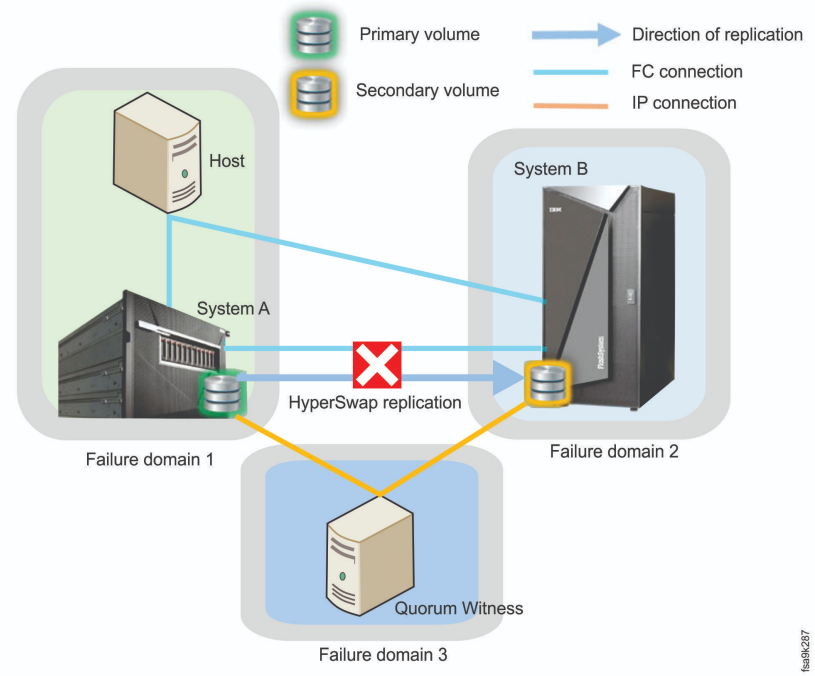

System A notifies the Quorum Witness about its taking sole responsibility for the relation. System B detects through the Quorum Witness that System A has taken sole responsibility for I/O, and stops serving I/O. It indicates to the host its inability to serve I/O by changing its port group state to *Unavailable*. At this point System B does not contain a synchronized image of the data and therefore is not ready for automatic failover.

As soon as connectivity is restored, the HyperSwap relation automatically re-synchronizes the peer volumes. System B changes the state of its port groups to *active-preferred*, and operates in redirect mode, meaning that it proxies all its I/O to System A. Once the relation is synced, readiness for automatic failover is resumed, and the Secondary volume resumes serving read requests directly.

System A - Quorum Witness connectivity failure

If connectivity between System A and Quorum Witness is disrupted, System A will periodically attempt to recover connectivity. While Quorum Witness connectivity is not available, System B remains ready for automatic failover. Automatic failover will be triggered, if connectivity between the systems A and B fails as well.

Figure 1. System A - Quorum Witness connectivity failure

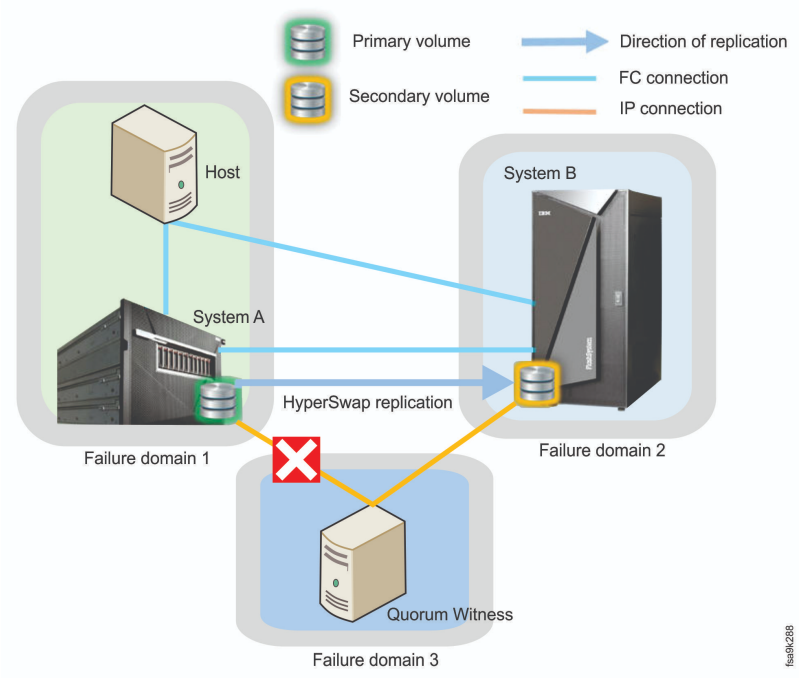

If the problem is not resolved, an additional failure of the peer connectivity occurs, and automatic failover becomes impossible.

As soon as connectivity is recovered, System A will send a heartbeat to the Quorum Witness. System B continues performing as usual.

System A - System B and System A - Quorum Witness connectivity failure or System A failure

If connectivity between System A and the Quorum Witness is disrupted, and connectivity between System A and System B is down, System A stops serving I/O and assumes that System B performed an automatic failover.

Figure 1. System A - System B and System A - Quorum Witness connectivity failure or System A failure

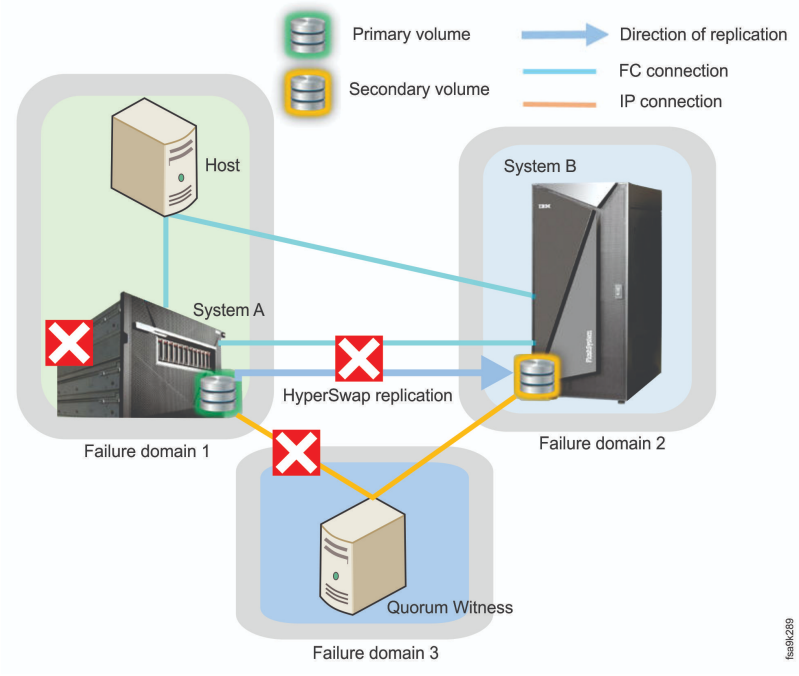

System B indeed triggers automatic failover by taking ownership of the Primary volume, that is by changing the role of the local volume from Secondary to Primary.

System A port group state changes to *Unavailable*. Automatic failover readiness is disabled, because only one copy of the data exists (on System B).

This scenario requires manual recovery. System A role must be changed to Secondary, and the HyperSwap relation needs to be activated. At this point data synchronization will begin. After the data is synchronized, the roles can be switched in order to restore the original configuration.

System B - Quorum Witness connectivity failure

If connectivity between System B and Quorum Witness is disrupted, System A will continue performing as usual. System B will periodically attempt to recover connectivity with the Quorum Witness, but readiness for automatic failover will be disabled.

Figure 1. System B - Quorum Witness connectivity failure

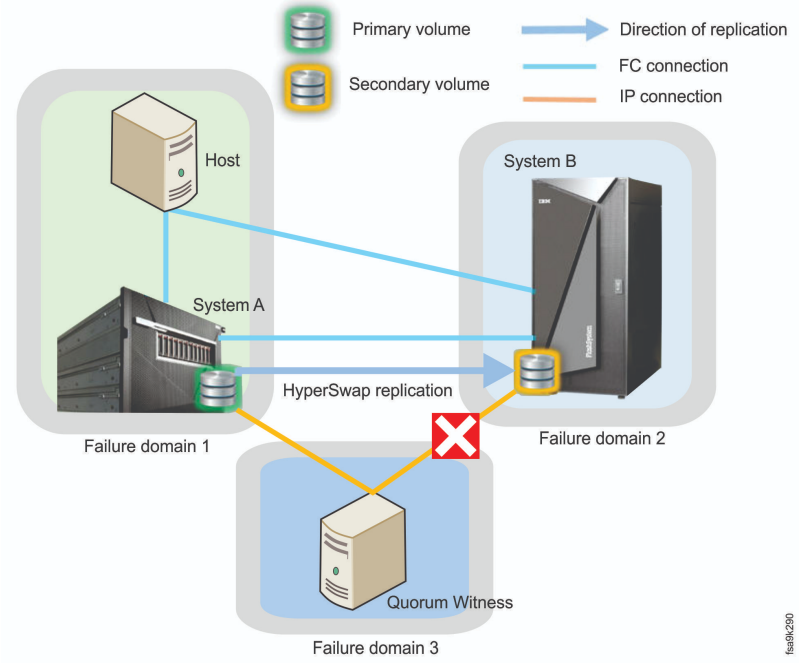

When Quorum Witness connectivity is restored, readiness for automatic failover is immediately and automatically resumed.

System B - System A and System B - Quorum Witness connectivity failure or System B failure

If connectivity between System B and System A is disrupted, the replication is disrupted as well, and the Secondary volume goes out of sync. System A will then assume sole responsibility for serving I/O. Readiness for automatic failover will be disabled.

Figure 1. System B - System A and System B - Quorum Witness connectivity failure or System B failure

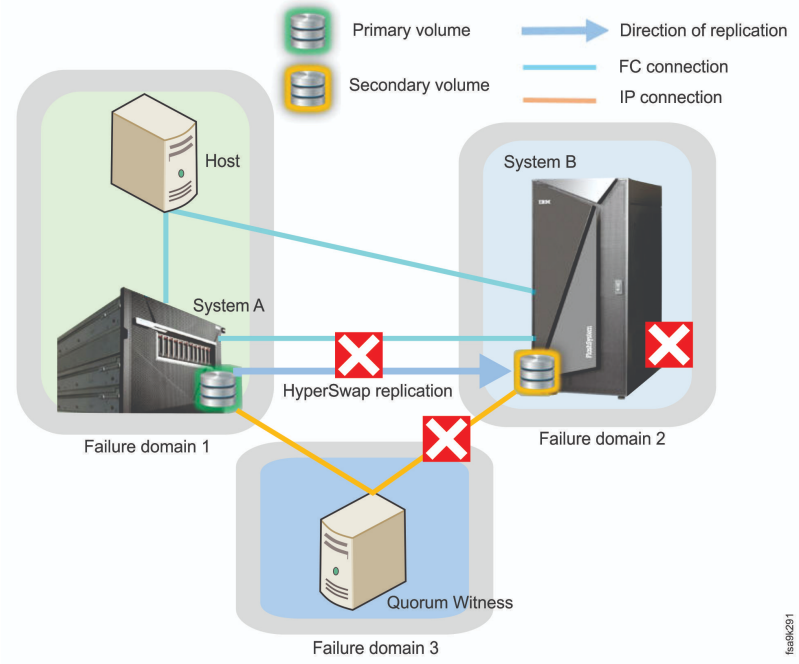

As soon as connectivity between System A and System B is restored, the volumes are automatically re-synchronized. System B changes the state of its port groups to *active/non-preferred*, and operates in redirect mode, meaning that it proxies any I/O to System A.

Once synchronization is complete and connectivity between System B and the Quorum Witness is restored, System B enables the HyperSwap relation and resumes handling read requests directly.

Automatic failover readiness is resumed once the data is synchronized and system B has connectivity to the Quorum Witness.

System A - Quorum Witness and System B - Quorum Witness connectivity failure

This scenario assumes that neither System A nor System B is connected to the Quorum Witness, while the System A - System B connectivity is in proper order. Automatic failover readiness will be disabled. System B, however, continues serving I/O, and data on both systems remains synchronized.

Figure 1. System A - Quorum Witness and System B - Quorum Witness connectivity failure

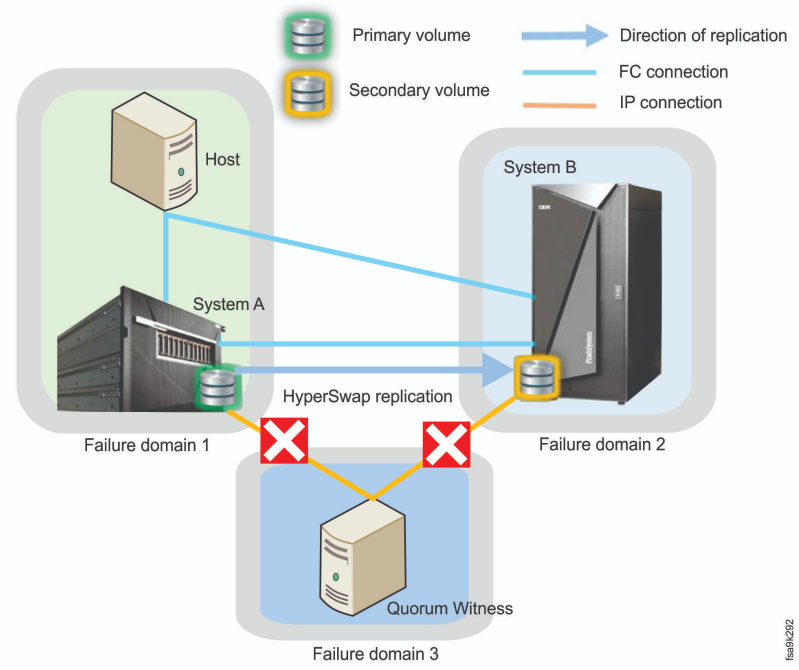

As soon as connectivity with the Quorum Witness is restored, both peers will automatically resume sending heartbeat messages, and the automatic failover will be automatically resumed.

System A I/O serving failure

If System A is functional, but is unable to serve I/O for any reason, it notifies System B of the failure, locks its HyperSwap volume, and sets the state of its port groups to *Unavailable*, thus blocking even I/O in flight. System B then automatically takes over the Primary volume and starts serving I/O.

Figure 1. System A I/O serving failure

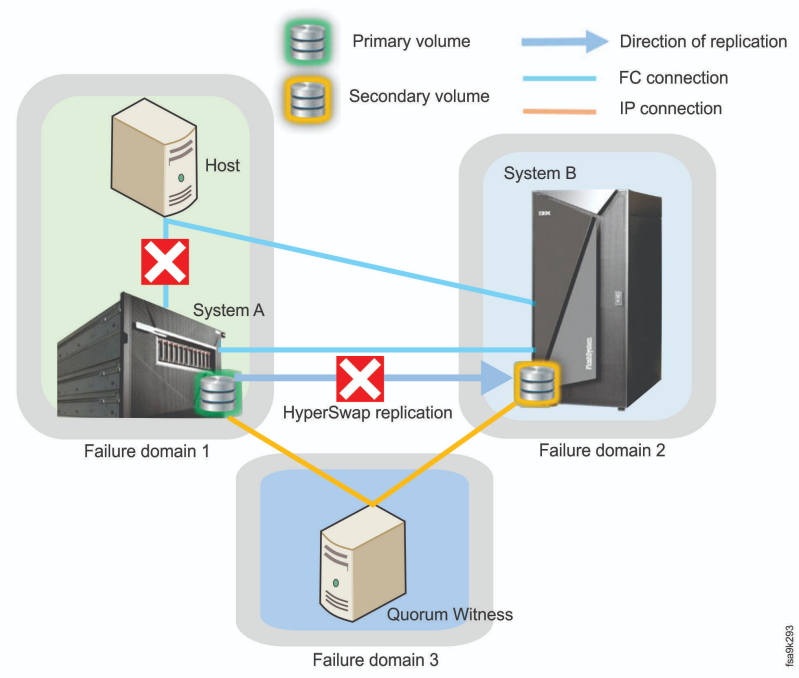

This scenario requires manual recovery. System A role must be changed to Secondary, and the HyperSwap relation needs to be activated. At this point data synchronization will begin. After the data is synchronized, the roles can be switched in order to restore the original configuration.

Total connectivity failure

This scenario is very unlikely and practically impossible, if the storage arrays and Quorum Witness are located in separate sites.

If all connectivity between System A, System B, and Quorum Witness is disrupted, none of the peers is aware of its counterpart's state. Therefore, both assume that the counterpart has taken over the volume, and stop serving I/O. The split-brain scenario is thus avoided, and data access will be completely blocked to avoid any data inconsistency.

This scenario requires manual recovery. Peer roles need to be reassigned as necessary, and automatic failover needs to be established.

Multi-site high availability and disaster recovery

The Multi-site high availability (HA) and disaster recovery (DR) feature allows customers to deploy field-proven high availability and disaster recovery solutions over multiple sites, keeping three concurrent copies of data.

As shown in Figure 1, a Multi-site HA/DR configuration is composed of three IBM FlashSystem A9000 or A9000R systems in a Multi-site star topology. To maximize the protection, the three systems should be deployed in three separate failure domains, with Quorum Witness deployed in System C's failure domains (see Failure domain for details).

A Multi-site relation comprises:

- One HyperSwap relation between the primary and secondary volumes on systems A and B
- One active asynchronous mirroring relation between the primary and tertiary volumes on systems A and C
- Optionally, one standby asynchronous mirroring relation between the secondary and tertiary volumes on systems B and C

This configuration provides HyperSwap active-active high availability, while keeping data mirrored to a third copy for disaster recovery, ensuring two levels of business continuity.

Figure 1. Multi-site HA/DR diagram

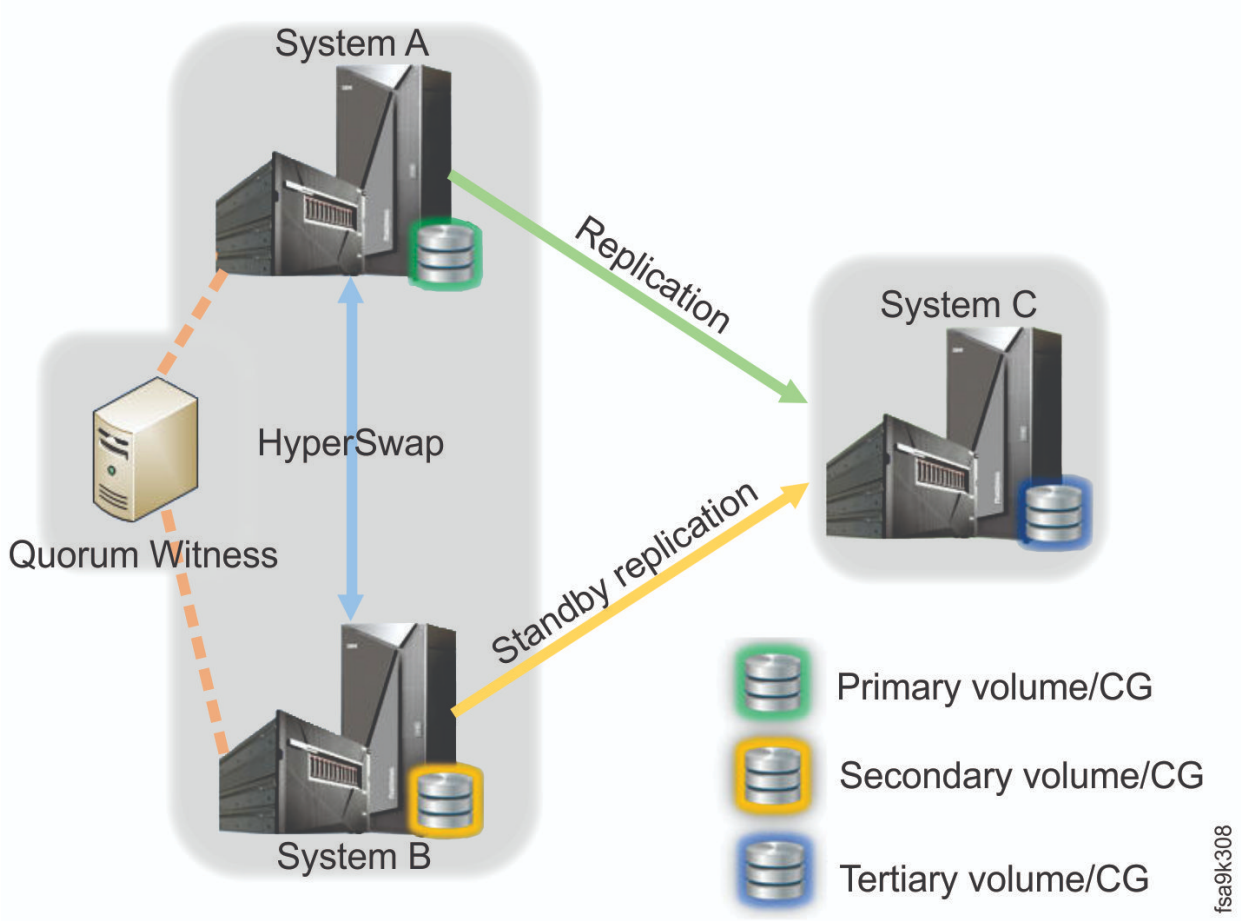

The Multi-site HA/DR feature is able to operate with volumes and consistency groups (CG). Any storage system in a Multi-site relation can have multiple Multi-site relations with volumes and consistency groups in different roles. For example, a system can have a volume or consistency group in the Primary role, and can have another volume or consistency group in the Secondary role, and a third volume or consistency group in the Tertiary role. Each of these are in different Multi-site relations.

- **Multi-site HA/DR advantages**
  As implemented in IBM FlashSystem A9000 and A9000R, Multi-site HA/DR offers flexibility, non-disruptive performance, and the ease of use.
- **Multi-site HA/DR architecture**
  A multi-site configuration is defined by extending a HyperSwap relation or an asynchronous mirror. In other words, defining the multi-site relation assumes that at least one fully initialized two-way HyperSwap or asynchronous mirroring relation exists.
- **Multi-site relation assembly**
  A multi-site relation can be assembled using Hyper-Scale Manager or CLI commands.
- **Multi-site HA/DR operations**
  All the operations related to Multi-site HA/DR configuration and management can be performed using IBM Hyper-Scale Manager (HSM) or the Command-Line Interface (CLI).
- **Recovery scenarios**
  IBM FlashSystem A9000 and IBM FlashSystem A9000R Multi-site HA/DR solution is designed to maintain non-disruptive host data access and recover data in case of a disaster that impacts both Systems A and B.

Multi-site HA/DR advantages

As implemented in IBM FlashSystem A9000 and A9000R, Multi-site HA/DR offers flexibility, non-disruptive performance, and the ease of use.

Any of the three peer systems participating in a multi-site relation can be either IBM FlashSystem A9000 or A9000R. System models do not have to be identical. In other words, a multi-site relation peer system can be model 415 or 425.

The Multi-site HA/DR feature does not require separate licensing.

The Multi-site HA/DR solution is designed to maintain non-disruptive host data access and recover data in case of a disaster. The administrator assembles it from existing relations. Any existing two-way relation (HyperSwap or asynchronous) can be extended to a multi-site relation without disruption while maintaining the data consistency of the existing relation.

Multi-site HA/DR is easy to configure and monitor, because the states of all the relations can be monitored though the multi-site relation. IBM Hyper-Scale Manager provides special management for the Multi-site HA/DR solution by allowing configuration and monitoring of the underlying relations via the parent multi-site relation. Special XCLI commands also allow for easy configuration and monitoring.

Multi-site HA/DR architecture

A multi-site configuration is defined by extending a HyperSwap relation or an asynchronous mirror. In other words, defining the multi-site relation assumes that at least one fully initialized two-way HyperSwap or asynchronous mirroring relation exists.

The multi-site relation runs on all systems (see Figure 1), with Primary and Secondary peers in a HyperSwap relation (see Peers in a Multi-site relation). One volume or configuration group has the Primary role for the HyperSwap relation or asynchronous mirror. Another volume or configuration group has the secondary role for the HyperSwap relation or asynchronous mirror. A third volume or configuration group has the tertiary role for the HyperSwap relation or asynchronous mirror.

Host connectivity and the asynchronous replication connectivity are supported over Fibre Channel (FC) or iSCSI links. The HyperSwap replication connectivity is supported over FC only.

- **Peers in a Multi-site relation**
  The Multi-site HA/DR technology introduces some new terms, in addition to those mentioned in the HyperSwap and asynchronous mirroring chapters.
- **Standby asynchronous mirror**
  The asynchronous mirroring relation between the secondary and tertiary volumes in a multi-site relation can be defined either in advance, at the time of the Multi-site relation assembly, or when needed for data recovery.
- **Snapshots in Multi-site HA/DR**
  Application consistent mirrored snapshots in the Multi-site HA/DR solution are supported in the context of the underlying HyperSwap and asynchronous relations. This means that mirrored snapshots on HyperSwap peers or on asynchronous mirror peers can be created.

Peers in a Multi-site relation

The Multi-site HA/DR technology introduces some new terms, in addition to those mentioned in the HyperSwap and asynchronous mirroring chapters.

In a regular, properly operating multi-site relation, depicted in Figure 1, the volume/CG roles are defined as follows.

Primary volume
: The primary volume in a HyperSwap relation with the secondary volume, and in an asynchronous mirroring relation with the tertiary volume.

Secondary volume
: The secondary volume in a HyperSwap relation between A and B. This volume also acts as the primary volume in the optional standby asynchronous mirroring relation between B and C.

Tertiary volume
: The secondary volume in two asynchronous mirroring relations, between A and C (active) and B and C (standby).

Figure 1. Multi-site HA/DR diagram

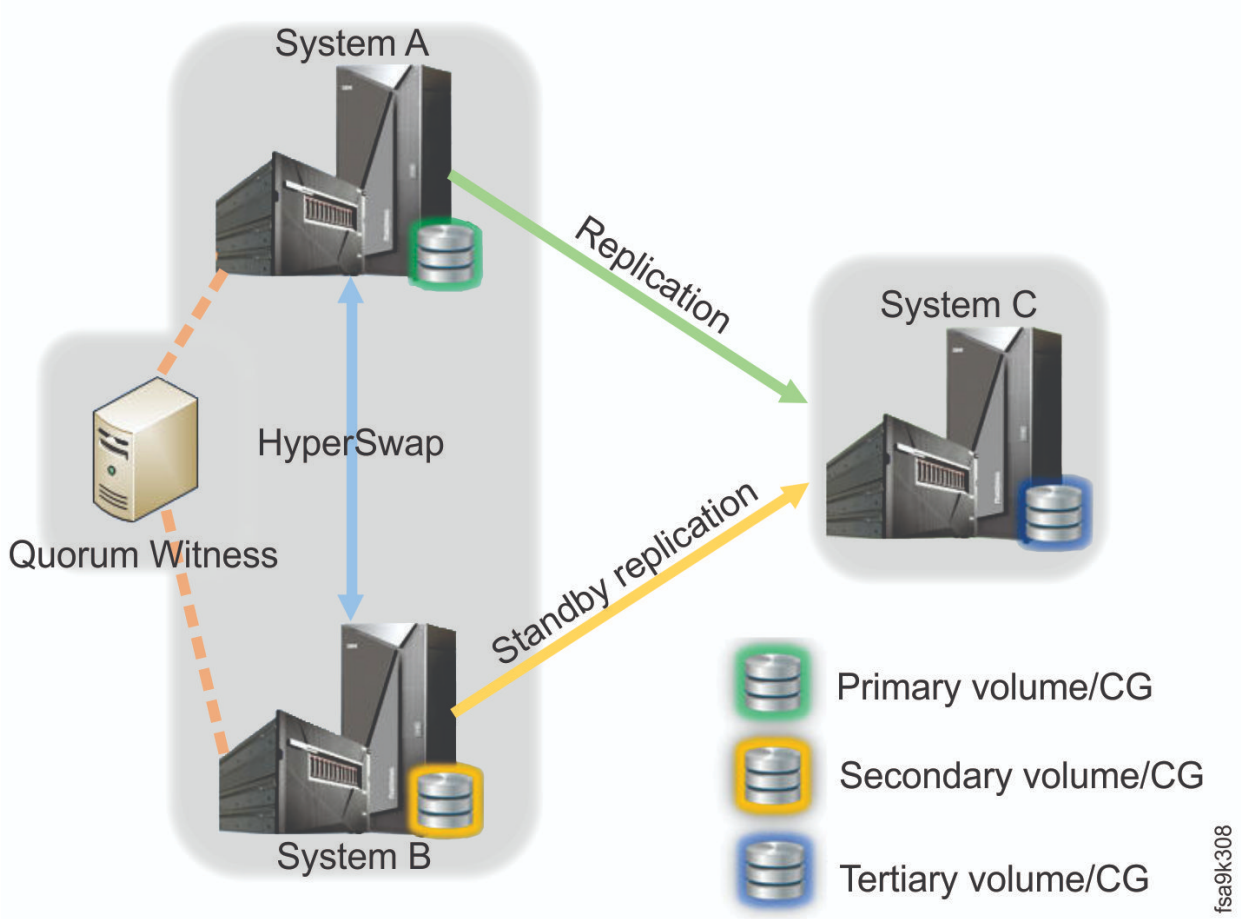

Important: As a counterpart to "Primary", "Secondary", and "Tertiary", the input and output syntax of CLI commands uses the legacy terminology of "Master", "sMaster", and "Slave". This inconsistency is a necessary compromise, required to avoid changes to older CLI commands that are in customer use, and also to keep the CLI terminology consistent across the board.
Disconnecting or losing the target of one of the relations does not affect the other relations in a multi-site configuration.

Standby asynchronous mirror

The asynchronous mirroring relation between the secondary and tertiary volumes in a multi-site relation can be defined either in advance, at the time of the Multi-site relation assembly, or when needed for data recovery.

When configuring the Multi-site relation using the XCLI, it is strongly recommended that the standby asynchronous mirror be configured at the time of the Multi-site assembly. If the relation is configured and registered, then in case there is a high availability event and a HyperSwap automatic failover takes place, disaster recovery continues automatically through the standby asynchronous mirror, and disaster recovery concurrency is maintained. This is optional in the XCLI because the XCLI allows for a more granular configuration.

When configuring Multi-site relation in HyperSwap Manager, the standby asynchronous mirror is configured automatically.

Defining the standby mirror in advance requires that the target connectivity between the secondary and tertiary volumes be in place between all systems when the multi-site relation is assembled. When defined, the standby mirror remains on standby under normal conditions and only becomes active automatically when System A fails, making the Primary volume inaccessible and triggering a failover from A to B in the A-B HyperSwap relation, as shown in Figure 1.

Figure 1. Standby replication after a HyperSwap failover

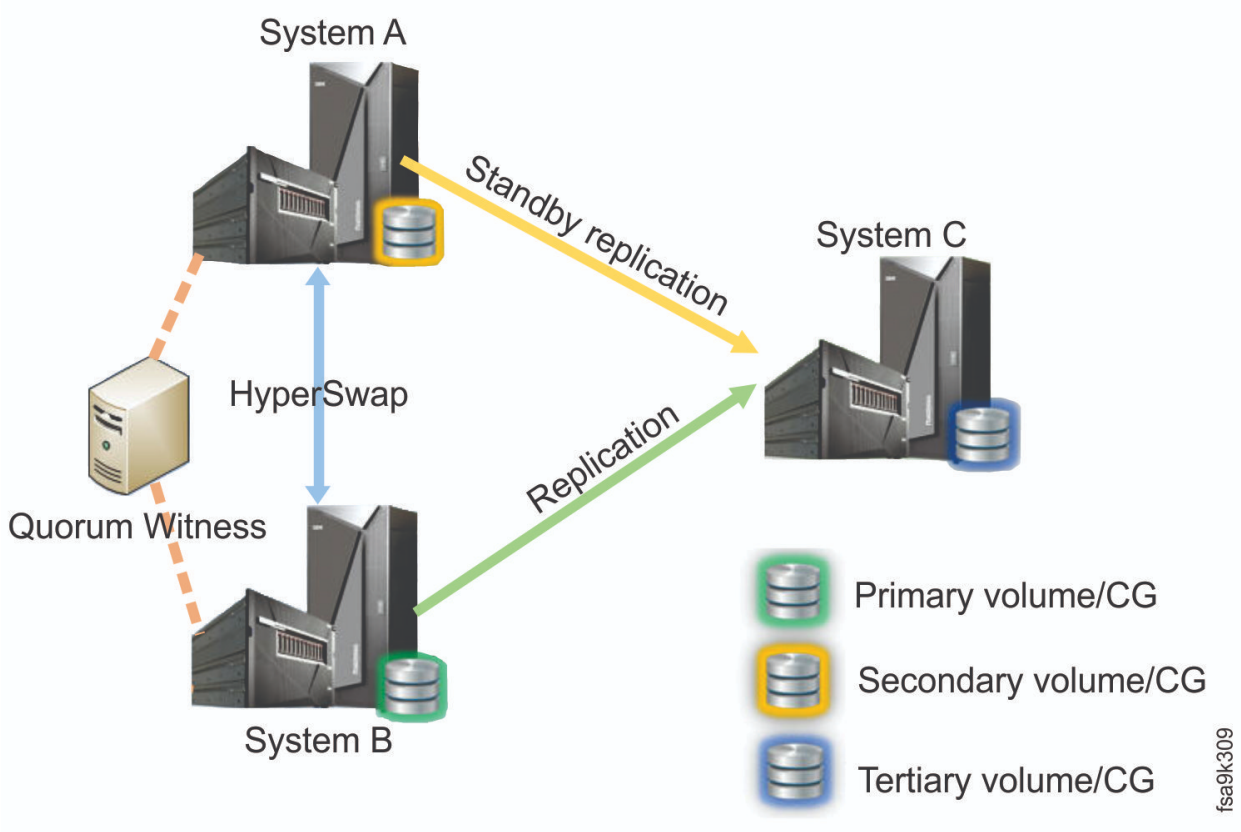

Snapshots in Multi-site HA/DR

Application consistent mirrored snapshots in the Multi-site HA/DR solution are supported in the context of the underlying HyperSwap and asynchronous relations. This means that mirrored snapshots on HyperSwap peers or on asynchronous mirror peers can be created.

System B holds and maintains the automatically created and refreshed most recent snapshot (MRS) and the last replicated snapshot (LRS) (see Mirroring-special snapshots for details). Upon a HyperSwap failover, the asynchronous replication from A to C stops, and the standby asynchronous relation between B and C is automatically activated. At this point, the MRS and LRS residing on System B are used to automatically resume the asynchronous replication between the Secondary and Tertiary copies from exactly the same point where the A to C replication stopped.

The snapshot-based technology implemented in the asynchronous replication between Systems A and C is identical to the mechanism used in the regular asynchronous mirroring, described in detail in Snapshot-based technology in asynchronous mirroring.

Multi-site relation assembly

A multi-site relation can be assembled using Hyper-Scale Manager or CLI commands.

As a preliminary requirement to assembling a multi-site relation, the following two-way relations must be defined first:

- HyperSwap relation between the volumes that in the multi-site relation will assume the roles of Primary and Secondary and vice versa. The HyperSwap relation status can be either Active and Sync, or Inactive.
- Asynchronous mirroring relation between the volumes that in the multi-site relation will assume the roles of Primary and Tertiary and vice versa. The asynchronous mirroring relation status can be either Active and Sync, or Inactive.
- If the standby relation is included, the connectivity between the volumes that in the multi-site relation will assume the roles of Primary and Tertiary must be configured and up.

For a multi-site relation to be successfully assembled, some other preliminary requirements must be met:

- All three Systems (A, B, and C) must be up and connected with each other.
- Only one of the relations - either HyperSwap or asynchronous mirroring - is active.
- Either the HyperSwap or the asynchronous mirroring relation is defined as *part of Multi-site*.

When a multi-site relation is assembled, it will be recognized on all the systems involved. In a consistency group, the multi-site relation is automatically added to each member volume.

The assembly of a multi-site consistency group relation requires the following to be accomplished:

1. Creation of either a HyperSwap consistency group between Systems A and B.
2. Creation of a consistency group on System C.

3. Creation of an asynchronous mirror between the consistency groups on Systems A and C.
   Note: The creation of the groups in 1 and 3 can be done in any order, but the creation of a consistency group on System C must be done in the middle.
4. Addition of the volumes from the consistency group on System A to the consistency group on System C, and establish an asynchronous mirror between the volume members of the consistency group on System A to the volumes on System C. This is performed with a single CLI command.
5. Creation of a consistency group level multi-site relation. The volume level multi-site relation is created automatically.
6. Establishing an asynchronous mirror between the consistency groups on Systems B and C.
7. Establishing an asynchronous mirror to the volume on System C for each volume in the consistency group on System B.
8. Registration of a consistency group between System B and C. This validates the configuration and ensures that there are no configuration issues that can prevent the automatic activation of the B-C mirror following a HyperSwap failover.
9. Activation of the asynchronous mirror between Systems A and C is activated. This triggers synchronization between Systems A and C.

- **Asynchronous mirror activation in a multi-site relation**
  Activating an asynchronous mirror that is part of a multi-site relation involves additional data integrity checks.
- **Adding a standby asynchronous mirror**
  If there is no standby asynchronous mirror defined at the point of the multi-site relation creation, you can add it at a later time, even after a failover.

Asynchronous mirror activation in a multi-site relation

Activating an asynchronous mirror that is part of a multi-site relation involves additional data integrity checks.

Before activating the asynchronous mirror, the system makes sure that either the relation between System A and System C is active, or the relation between System B and System C is active, and not both at the same time. This is to make sure that old data does not overwrite new data. For this reason, activation on an asynchronous mirror with is part of a multi-site relation must use the multi-site level command, and not the mirror level command.

Adding a standby asynchronous mirror

If there is no standby asynchronous mirror defined at the point of the multi-site relation creation, you can add it at a later time, even after a failover.

This can be the case when the multi-site relation was created using the XCLI. When the multi-site relation is created using HyperSwap Manager, HyperSwap Manager automatically creates the multi-site relation using the best practices, including the standby asynchronous mirror.

The asynchronous mirror between B and C can start from the last consistent copy created by A, even if the B-C asynchronous mirror is added after a failover.

As a preliminary requirement, the Secondary-Tertiary asynchronous mirroring relation must be created first.

When a standby asynchronous mirror is successfully created, it must be registered in the multi-site relation. In a consistency group, the registration operation can be carried out at the consistency group level. As a result, the standby relation is registered on every member volume.

Multi-site HA/DR operations

All the operations related to Multi-site HA/DR configuration and management can be performed using IBM Hyper-Scale Manager (HSM) or the Command-Line Interface (CLI).

HSM simplifies management for the Multi-site HA/DR feature. This makes Multi-site HA/DR easy to configure and monitor.

Similar to the rest of IBM FlashSystem A9000 and A9000R features and capabilities, Multi-site HA/DR can be also configured and managed using CLI.

- **Monitoring a Multi-site relation**
  The current status of a Multi-site relation and the two-way relations it is assembled from can be monitored at any time using IBM Hyper-Scale Manager or a CLI command.
- **Changing roles in a Multi-site relation**
  The necessity to change the role of a Multi-site relation peer may arise after an automatic or manual failover or as part of the disaster recovery procedure.
- **Deleting a Multi-site relation**
  The Multi-site relation is represented in the storage system software by an explicit *Multi-site* object. Deleting the Multi-site relation is done by simply deleting the relevant object.
- **Downgrading a Multi-site relation to a two-way mirroring relation**
  A Multi-site relation can be reverted to a two-way relation, which can be necessary in a disaster recovery situation.

Monitoring a Multi-site relation

The current status of a Multi-site relation and the two-way relations it is assembled from can be monitored at any time using IBM Hyper-Scale Manager or a CLI command.

Multi-site HA/DR has multiple states, or conditions, of operation. While each individual two-way relation has its own state, the Multi-site relation has a global state too. The existing options are explained in Table 1.

Table 1. Global states of a Multi-site relation

Value	Meaning

Value	Meaning
`Inactive`	Both two-way relations that make part of a Multi-site relation are inactive.
`Initializing`	One or both two-way relations that make part of a Multi-site relation are in the *init* state.
`Operational`	The HyperSwap relation is synchronized, automatic failover is active, the asynchronous mirror is RPO OK, Secondary is ready to assume the ownership of the disaster recovery (DR).
`Degraded`	The asynchronous mirror is RPO lagging.

`Compromised`	This state can be triggered by one of the following:
	• The connection is disrupted between Primary and Secondary and/or Primary and Tertiary, or between Secondary and Tertiary. The latter is applicable only if a standby mirror is configured.
	• Resynchronization between Primary and Secondary and/or Primary and Tertiary is in progress.
	• The machine holding Secondary cannot assume the ownership of disaster recovery (DR), because Secondary does not have an updated MRS (most recent snapshot) or LRS (last replicated snapshot.)
	• A partial role change, either automatic or manual, has occurred between Primary and Secondary or Primary and Tertiary.
	• Secondary cannot perform automatic failover. This may occur because:
	  ◦ Secondary is not synchronized
	  ◦ Secondary is disconnected from the Quorum Witness
	  ◦ automatic failover is disabled by the user
	  ◦ connectivity between Primary and Secondary has failed

`Inconsistent`	A relation between Primary and Tertiary is not defined. This can happen when a standby relation was not created and Secondary becomes Primary.

The states of the registered standby mirror between Secondary and Tertiary are listed in Table 2. An asynchronous mirroring relation is considered a registered standby mirror if it is defined on a Multi-site Secondary volume or consistency group, and is inactive and had either been registered or active in the past. The existing options are explained in the table below.

Table 2. States of a standby mirror

Value	Meaning
`Up`	The standby mirror is defined, registered and connected.
`Down`	The standby mirror is defined and registered, but disconnected.
`N/A`	The standby mirror is either undefined or unregistered.

Changing roles in a Multi-site relation

The necessity to change the role of a Multi-site relation peer may arise after an automatic or manual failover or as part of the disaster recovery procedure.

The role of the local Multi-site relation peer may need to change in the following cases:

- When restoring the original configuration after an automatic or manual failover.
- When performing a manual failover when the automatic failover did not take place.
- When recovering from a DR (disaster recovery) site.

The following role changes are allowed:

- Primary to Secondary
  For a successful role change from Primary to Secondary, the volume can be in any phase, except Initializing. Upon changing its role, the Primary ceases serving host requests, and is set to act as Secondary and accept replication from the other peer.
- Secondary to Primary
  After Secondary is successfully changed to a Primary, the volume (or consistency group) starts accepting requests from hosts. Upon explicit activation, it starts replicating to the peer volume (the former Primary).
- Tertiary to Primary
  This role change is available only for a disconnected Tertiary, that is not connected to either Primary or Secondary. After such Tertiary becomes Primary, it cannot be incorporated into an existing multi-site relation. Instead, a new Multi-site relation must be defined.

Deleting a Multi-site relation

The Multi-site relation is represented in the storage system software by an explicit *Multi-site* object. Deleting the Multi-site relation is done by simply deleting the relevant object.

Before the *Multi-site* relation can be deleted, the active relations on Primary must be deactivated. Then, the deletion of the *Multi-site* object is done on System A. After the Multi-site relation is deleted, the HyperSwap asynchronous mirror relations remain. The user can delete or activate them.

If a Multi-site relation cannot be deleted on every peer due a system failure or connectivity problem, it can be deleted partially. This allows deleting a Multi-site configuration when the connectivity between the systems is down.

If a permanent failure has occurred, and the Multi-site configuration cannot be recovered, the Multi-site relation can be removed from the local site, even though the peers may be disconnected.

When a Multi-site relation is deleted, the Multi-site internal snapshots on Secondary become external and available for the user to manage. If a new Multi-site relation is created for the same volume or consistency group, those snapshots must be deleted.

Downgrading a Multi-site relation to a two-way mirroring relation

A Multi-site relation can be reverted to a two-way relation, which can be necessary in a disaster recovery situation.

The downgrade operation chooses which two-way relation should be kept. The other two-way relations constituting the Multi-site relation get deleted automatically (when using Hyper-Scale Manager) or go into a standby state (when using the CLI command). They can be later deleted individually or reused to create another Multi-site relation.

Recovery scenarios

IBM FlashSystem A9000 and IBM FlashSystem A9000R Multi-site HA/DR solution is designed to maintain non-disruptive host data access and recover data in case of a disaster that impacts both Systems A and B.

For detailed information about recovery scenarios, refer to the IBM HyperSwap and Multi-site HA/DR for IBM FlashSystem A9000 and A9000R IBM Redbooks publication.

Migrating data

The data migration feature enables the transfer of large amounts of data with continuous data access during the process.

The use of any new storage system frequently requires the transfer of large amounts of data from another storage system. The data migration feature enables continuous data access during the actual data transfer, thus reducing the required downtime time to the minimum, sometimes from days or hours to minutes. In addition, data migration is applicable to any source block storage.

The following table summarizes the best practice migration methods.

Table 1. Best practice migration methods

The storage system from which to migrate data	Best practice migration method
Non-IBM® storage system	Regular data migration
Any IBM storage system, except IBM XIV® Gen3, IBM FlashSystem® A9000, and FlashSystem A9000R	Regular data migration
IBM XIV Gen3, versions prior to 11.6.2.a	Regular data migration
IBM XIV Gen3, version 11.6.2.a or later	Hyper-Scale Mobility
IBM FlashSystem A9000R, versions prior to 12.2.0	Regular data migration
IBM FlashSystem A9000R version 12.2.0 or later	Hyper-Scale Mobility
IBM FlashSystem A9000	Hyper-Scale Mobility

Note: When the origin system is IBM XIV Gen3 (version 11.6.2.a or later), IBM FlashSystem A9000 (any version), or FlashSystem A9000R (version 12.2.0 or later), the best practice migration method is IBM Hyper-Scale Mobility, since it is non-disruptive. For more information, see Volume migration with IBM Hyper-Scale Mobility.

- **Data migration overview**
- **I/O handling in data migration**
- **Data migration stages**
- **Handling failures**

Data migration overview

The data migration feature enables the smooth transition of a volume from any block storage to IBM FlashSystem® A9000 or A9000R, by:

- Connecting the associated host to IBM FlashSystem A9000 or A9000R and providing the host with direct access to the most up-to-date data even before data has been copied from the origin storage system.
- Transparently copying the contents of the origin storage system to IBM FlashSystem A9000 or A9000R as a background process.

During data migration, the host is connected directly to IBM FlashSystem A9000 or A9000R and is disconnected from the origin storage system to avoid a data integrity issue. IBM FlashSystem A9000 or A9000R is connected to the origin storage system, as shown in the Data migration topology chart below.

Figure 1. Data migration topology

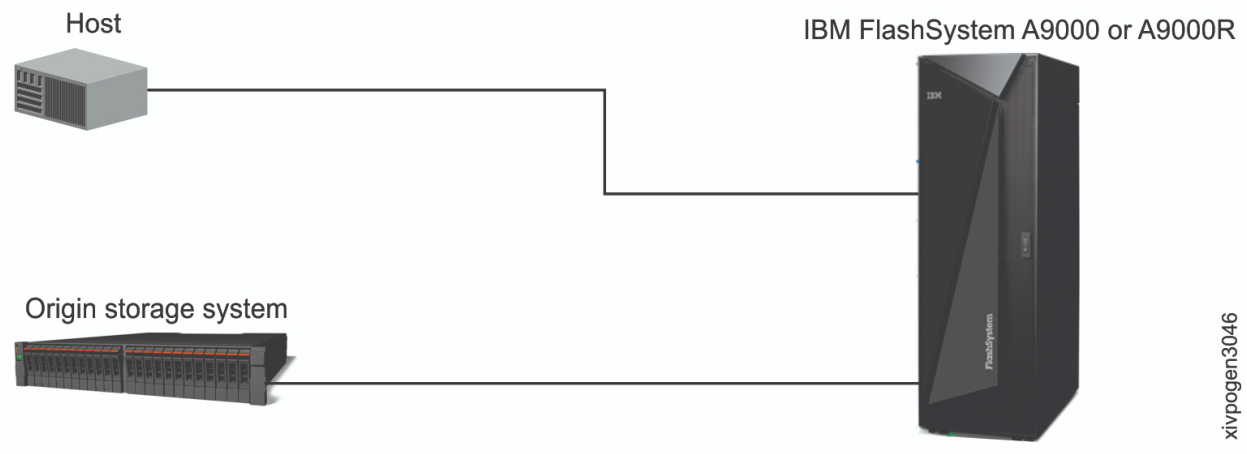

The communication between the host and IBM FlashSystem A9000 or A9000R and the communication between IBM FlashSystem A9000 or A9000R and the origin storage system is over a Fibre Channel or iSCSI connection.
Important: SCSI Hardware accelerated locking (ATS command) is not supported for data migration. When host operating systems that use ATS are involved, customers should either disable ATS usage on these hosts for the duration of the migration, or consider other migration methods.
Customers who can use VMware storage vMotion for the required volume migration, should prefer this method for two reasons:

- Storage vMotion does not disrupt the involved VMs.
- While valid, disabling ATS (for VAAI-capable hosts) has operational ramifications as described in the VMware Knowledge Base.

I/O handling in data migration

Serving read requests

During this stage, IBM FlashSystem® A9000 or A9000R will serve all the host's data read requests. This will be performed in a transparent manner without requiring any action by the host, as follows:

- If the requested data has already been copied to IBM FlashSystem A9000 or A9000R, it is served from IBM FlashSystem A9000 or A9000R.
- If the requested data has not yet been copied, IBM FlashSystem A9000 or A9000R will retrieve it from the origin storage system and then serve it to the host.

Serving write requests

During this stage IBM FlashSystem A9000 or A9000R will serve all host's data write requests. This will be performed in a transparent manner without requiring any action by the host.

Data migration provides the following two alternative configurations for handling write requests from a host:

Source updating:
A host's write requests are written by IBM FlashSystem A9000 or A9000R to IBM FlashSystem A9000 or A9000R, as well as to the origin storage system. In this case, the origin storage system remains completely updated during the background copying process. Throughout the process, the volume of the origin storage system maintains an up-to-date copy of the data.
Write commands are performed synchronously, so IBM FlashSystem A9000 or A9000R only acknowledges the write operation after writing to itself, writing to the origin storage system, and receiving an acknowledgment from the origin storage system. Furthermore, if, due to a communication error or any other error, the writing to the origin storage system fails, IBM FlashSystem A9000 or A9000R will report to the host that the write operation has failed.

Source updating is the recommended method for handling write requests from the host, because it provides the ability to fall back if the data migration process cannot be successfully completed.

No source updating
A host's write requests are only written by IBM FlashSystem A9000 or A9000R to IBM FlashSystem A9000 or A9000R and are not written to the origin storage system. In this case, IBM FlashSystem A9000 or A9000R is not updated during the background copying process. Therefore, the two storage systems will never be identical once a write was done by the host to IBM FlashSystem A9000 or A9000R. The volume of the origin storage system will remain intact and will not be changed throughout the data migration process.

Data migration stages

Figure 1 depicts the process of migrating a volume from a previous storage system to the new storage system. It also shows how the new storage system synchronizes its data with the previous storage system, and how it handles the data requests of a host throughout all these stages of synchronization.

Figure 1. Data migration steps

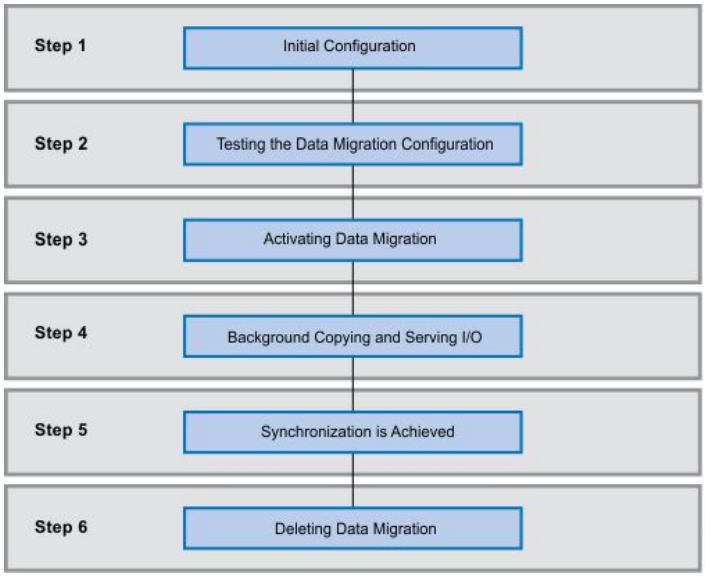

Initial configuration

The following prerequisite actions must be performed prior to initiating data migration:

- Format the volume on IBM FlashSystem® A9000 or A9000R.
- Disconnect the actual host from the origin storage system (see Testing the data migration configuration for details).
- Connect IBM FlashSystem A9000 or A9000R to the origin storage system whose data it will be serving. Because the origin storage system perceives IBM FlashSystem A9000 or A9000R as a host, IBM FlashSystem A9000 or A9000R must connect to the origin storage system as a SCSI initiator.
- Make sure that the volumes on the origin storage system and on IBM FlashSystem A9000 or A9000R have an equal number of 512 byte blocks. This is verified upon activation of the data migration process.

You can then initiate data migration and configure all hosts to work directly and solely with IBM FlashSystem A9000 or A9000R.

Testing the data migration configuration

Before connecting the host to IBM FlashSystem A9000 or A9000R, test the data migration definitions to verify that IBM FlashSystem A9000 or A9000R can access the origin storage system.

As a mandatory prerequisite for this stage, the actual host must be disconnected from IBM FlashSystem A9000 or A9000R (see Initial configuration). Alternatively, one test can be performed while the host is still connected, and another test can be performed after disconnecting the host. This allows the customer to ensure that the original volume size has not changed.

Activating data migration

After you have tested the connection between IBM FlashSystem A9000 or A9000R and the origin storage system, activate data migration and connect the host to IBM FlashSystem A9000 or A9000R. From this point forward, the host reads and writes data to IBM FlashSystem A9000 or A9000R, and IBM FlashSystem A9000 or A9000R will read and optionally write to the origin storage system.

Once activated, deactivating data migration will cause host I/O to fail. Use deactivation only for recovery after a failure, if the data migration configuration cannot be recovered.

Background copying and serving I/O operations

Once data migration is initiated, it will start a background process of sequentially copying all the data from the origin storage system to IBM FlashSystem A9000 or A9000R.

Synchronization is achieved

After all of a volume's data has been copied, the data migration achieves synchronization. After synchronization is achieved, all read requests are served from IBM FlashSystem A9000 or A9000R.

If source updating is enabled, IBM FlashSystem A9000 or A9000R will continue to write new data to both itself and the origin storage system until data migration settings are deleted (for details on source updating, see Serving write requests.)

Deleting data migration

When FlashSystem A9000 or A9000R is synchronized with the origin storage system, the data migration process can be safely deactivated and deleted together with the source volume.

If the data migration configuration is incorrect (for example, a wrong source volume is selected or the defined data block size is unsuitable) or if data migration is activated merely for a proof of concept, the process can be forcefully terminated, deactivated, and deleted. In this case, even if the target volume may seem

redundant, make sure that any written data is preserved before deleting it. This is especially important if source updating was not activated during data migration, because the target volume may contain data that is missing from the source volume.

If source updating was activated during data migration, all written data will remain on the source volume.

Handling failures

Upon a communication error or the failure of the origin storage system, I/Os received on IBM FlashSystem® A9000 or A9000R are likely to fail as source updating writes cannot be sent to the origin storage system, and data not yet copied to IBM FlashSystem A9000 or A9000R cannot be read.

If IBM FlashSystem A9000 or A9000R encounters a Medium Error on the origin storage system (meaning that IBM FlashSystem A9000 or A9000R cannot read a data block on the origin storage system), then the state of the same data block in IBM FlashSystem A9000 or A9000R will indicate a Medium Error, even though the backend persistent media in IBM FlashSystem A9000 or A9000R have not failed.

Volume migration with IBM Hyper-Scale Mobility

IBM® Hyper-Scale Mobility enables a non-disruptive migration of volumes from an IBM FlashSystem® A9000 or A9000R to another IBM FlashSystem A9000 or A9000R storage system, and from XIV® Gen3 to an IBM FlashSystem A9000 or A9000R.

IBM Hyper-Scale Mobility dramatically minimizes time, complexity, risk, cost and variety of tools required to non-disruptively vacate complete storage arrays, that serve different host operating systems and applications. Typically, it helps achieve data migration in the following scenarios:

- Migrating data out of an over-provisioned system.
- Migrating all the data from a system that will be decommissioned or re-purposed.
- Migrating data to another storage system to achieve adequate (lower or higher) performance, or to load-balance systems to ensure uniform performance.
- Migrating data to another storage system to load-balance capacity utilization.

Note: Starting from version 12.2.1, customers can non-disruptively move applications out of multiple XIV Gen3 storage systems and them consolidate them onto a single IBM FlashSystem A9000 or IBM FlashSystem A9000R.

Figure 1. Cross-generation volume migration

For support information on IBM Hyper-Scale Mobility, see the feature availability matrix.

Important: When using IBM Hyper-Scale Mobility on Storage Area Network (SAN) Boot Volumes, the Proxy targets need to be replaced with the Owner targets on the adapter BIOS. In some cases, this can be done online from the operating system, by using vendor's online tools for x86 servers (for example, Qlogic QConvergeConsole), or by using specific commands to edit the boot paths order for Unix servers. In other cases, manual changes of the adapter BIOS are required. When migrating volumes from a storage system that does not support IBM Hyper-Scale Mobility, a best practice is to use the regular data migration. It minimizes the duration of the disruption by concurrently allowing both data access and data migration.

- **Volume migration process**
  IBM Hyper-Scale Mobility moves a volume from one system to another, while the host keeps using the volume.

Volume migration process

IBM® Hyper-Scale Mobility moves a volume from one system to another, while the host keeps using the volume.

To accomplish this, I/O paths are manipulated by the storage, without involving host configuration, and the volume identity is cloned on the target system. In addition, direct paths from the host to the target system need to be established, and paths to the original host can finally be removed. Host I/Os are not interrupted throughout the migration process.

The key stages of the IBM Hyper-Scale Mobility and the respective states of volumes are shown in Figure 1 and explained in detail in Table 1.

For an in-depth practical guide to using IBM Hyper-Scale Mobility, refer to the following Redpaper publication: IBM Hyper-Scale Mobility Overview and Usage.

Figure 1. Volume migration flow of the IBM Hyper-Scale Mobility

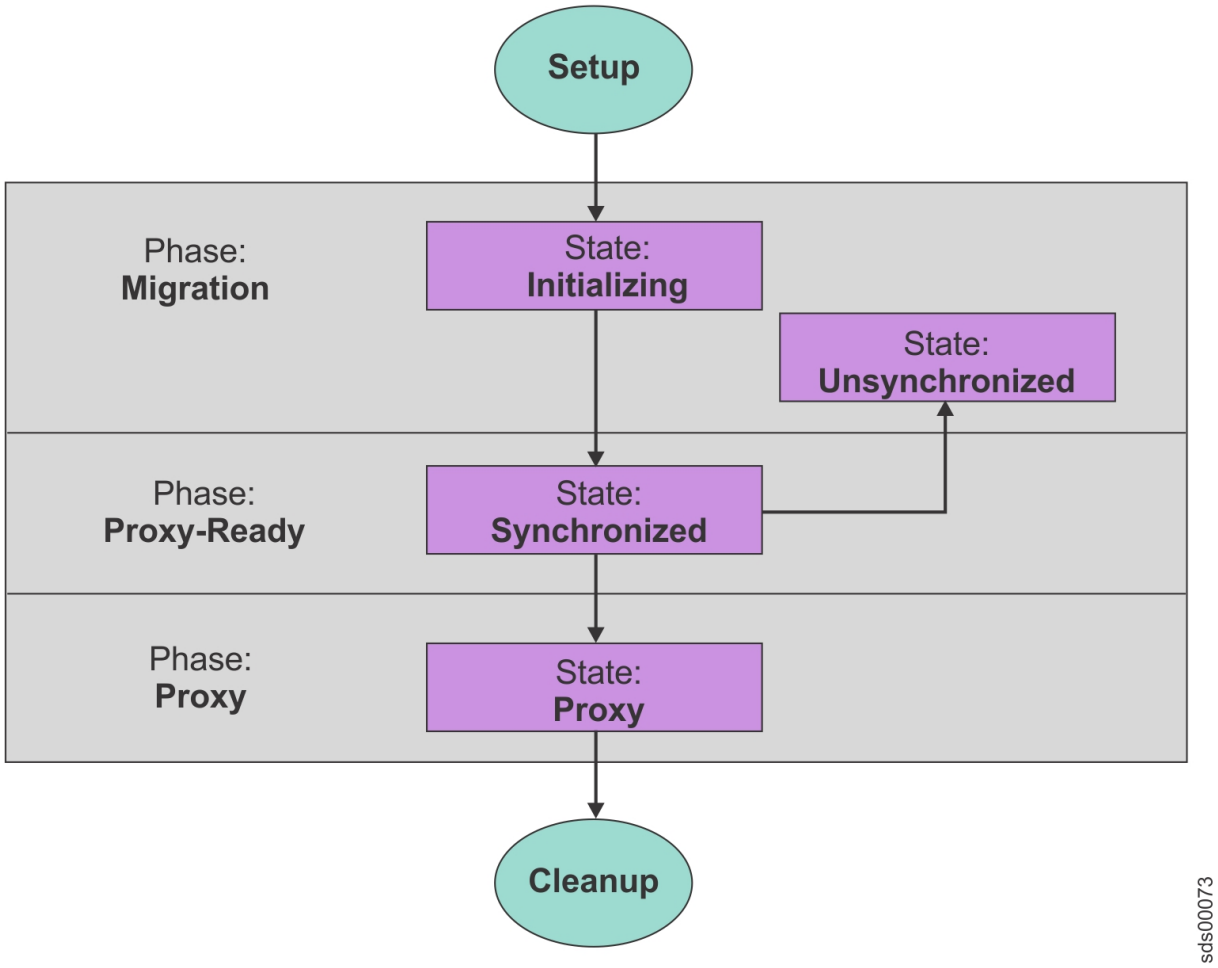

Table 1. Volume migration stages

Stage	Description	Source and destination volume states
Setup	A volume is automatically created at the destination storage system with the same name as the source volume. The relation between the source and destination volumes is established.	The two volumes are not yet synchronized.
Migration	New data is written to the source and replicated to the destination.	**Initializing** - The content of the source volume is copied to the destination volume. The two volumes are not yet synchronized. This state is similar to the Initialization state of synchronous mirroring. As long as the source instance cannot confirm that all of the writes were acknowledged by the destination volume, the state remains Initializing.
Proxy-Ready	The replication of the source volume data is complete when the destination is synchronized. The source serves host writes as a proxy between the host and the destination. The system administrator issues a command that moves the volume migration relation to the proxy. Next, the system administrator maps the host to the destination. In this state, a single copy of the data exists on the destination and any I/O directed to the source is redirected to the destination.	**Synchronized** - The source was fully copied to the destination. This state is similar to the Synchronized state of synchronous mirroring.

Stage	Description	Source and destination volume states
Proxy	New data in written to the source and is migrated to the destination. The proxy serves host requests as if it were the target, but it actually impersonates the target. Note: It is a best practice for the IBM Hyper-Scale Mobility to remain in Proxy state for a minimal period.	**Proxy -** The source acts as a proxy to the destination.
Cleanup	After validating that the host has connectivity to the destination volume through the new paths, the storage administrator unmaps the source volume on the source storage system from the host. Then the storage administrator ends the proxy and deletes the relation.	

Mirroring Resources

Synchronous mirroring, asynchronous mirroring, Hyper-Scale Mobility, HyperSwap, Multi-site High Availability/Disaster Recovery (HA/DR) and data migration, are all based on IBM FlashSystem® A9000 and A9000R data replication technology.

As such, they consume from a single pool of internal data replication resources, called *couplings*. Couplings are consumed as follows:

- **Synchronous** replication of a volume or a consistency group (CG) consumes **1 coupling**.
- **Asynchronous** replication of a volume or a CG consumes **3 couplings**.
- Except in asynchronous mirroring, and the asynchronous leg of a Multi-site HA/DR, any data replication is synchronous.
- In a Multi-site HA/DR relation, the standby asynchronous relation does not consume couplings.
- When a consistency group is replicated, couplings are consumed by the CG itself, and by each of its volumes, as specified above.
- The system supports a maximum of 3072 couplings.
  Important: FlashSystem A9000R systems support a maximum of 3072 mirror couplings. These couplings are shared between all and any combinations of synchronous, asynchronous, HyperSwap and Multi-site relations. The storage system high availability feature allows for a maximum of 1536 HyperSwap pairs. These pairs are a part of the total allowance of 3072 mirror couplings that the system supports.
  For more information about HyperSwap, see High availability with HyperSwap.

Note: For more examples and more information, refer to the *IBM FlashSystem A9000 and A9000R Business Continuity Solutions* Redbooks publication.

- **Examples for replication couplings**

Examples for replication couplings

The following examples show how to plan the allocation of mirroring resources. Each of the examples is based on a system currently configured with exactly one **asynchronous** Consistency Group (CG) that contains 5 volumes.

Maximum number of synchronous mirrors

Question: On top of the existing replications, how many additional individual volumes can now be **synchronously** mirrored?

Answer: 3054

Explanation: In this case, the number of couplings that are already used is 18 :

18 = (**3 couplings** x 1 **asynchronous** CG) + (**3 couplings** x 5 volumes within the **asynchronous** CG)

The total number of couplings in the system is the maximum number of synchronous mirrors, which is 3072. Therefore, 3054 couplings are left:

3054 = 3072 – 18

Therefore, since a synchronously mirrored volume consumes 1 couplings, 3054 volumes can be synchronously mirrored individually (that is, not in a CG).

Maximum number of asynchronous mirrors

Question: On top of the existing replications, how many additional individual volumes can now be asynchronously mirrored?

Answer: 1018

Explanation: 3054 couplings are left. Each asynchronously mirrored volume requires 3 couplings. Therefore, 1018 volumes can be asynchronously mirrored.

1018 = 3054/3

Maximum number of CGs with 3 volumes using Multi-site HA/DR

Question: On top of the existing replications, how many CGs, each with 3 volumes can be protected using Multi-site HA/DR?

Answer: 190

Explanation: Every Multi-site HA/DR relation contains:

- 1 HyperSwap relation.
- 1 active asynchronous relation.
- 1 standby asynchronous relation.

Therefore, a single Multi-site HA/DR object (a volume or CG) consumes 4 couplings:

- 1 coupling for the HyperSwap relation
- 3 couplings for asynchronous relation
- 0 couplings for the standby asynchronous relation
- Total = 4 couplings

Therefore, a single Multi-site HA/DR consistency group with 3 volumes consume 16 couplings:

16 = 1 consistency group x 4 + 3 volumes x 4

Finally, the number of such CGs that can be created with the free couplings is 190:

190 = ⌊3054/16⌋

Note the use of floors to round down the result. In this example, 14 couplings are left unused, because 16 couplings are the minimum for a single Multi-site HA/DR CG with 3 volumes.

14 = 3054 – 190 x 16

Data-at-rest encryption

The storage system secures all written data with industry-standard AES-256 encryption for data-at-rest.

Data-at-rest encryption protects the grid controller SSD disks and flash enclosure MicroLatency® modules against exposure of sensitive data on discarded or stolen media by ensuring that the data cannot be read as long as the key used to encrypt the data is secured. Encryption is carried out at the hardware level to avoid any performance impact.

Each MicroLatency module has FPGA control for data-at-rest encryption. Data-at-rest is protected by an Advanced Encryption Standard (XTS-AES) algorithm using the 256-bit symmetric option in xor-encrypt-xor (XEX)-based tweaked-codebook mode with ciphertext stealing (XTS) mode, as defined in the IEEE1619-2007 standard.

In addition, an HMAC-SHA256 algorithm is used to create a hash message authentication code (HMAC) for corruption detection. A system-generated cyclic redundancy check (CRC) is applied for additional protection.

Encryption can be enabled during the installation of the system or at any time later. While encryption is not enabled, the system might not meet customers or legal compliance standards and the data might not be protected against security issues. Encryption can be disabled only when no volumes are defined.

For an extended discussion about data-at-rest encryption, refer to the Data-at-rest Encryption for the IBM Spectrum Accelerate Family Redbooks publication . For more information about digital certificates and certificate authorities, refer to the Managing Digital Certificates across the Enterprise Redbooks publication.

- **Encryption key management schemes**
  Encryption key management can be carried out through an external or an internal scheme.

Encryption key management schemes

Encryption key management can be carried out through an external or an internal scheme.

External encryption key management scheme

With an external encryption key management scheme, encryption keys are stored separately from the data, thereby presenting a secured and well-defined interface for key services.

The separation of key storage from data storage and key management is accomplished with IBM® Security Key Lifecycle Manager (SKLM), Thales KeySecure (Gemalto SafeNet KeySecure) server, or Thales CipherTrust Manager.
Note: Gemalto SafeNet KeySecure is now known as Thales KeySecure and will be referred to that way throughout the rest of this guide.
The supported versions of IBM Security Key Lifecycle Manager (SKLM) are 2.6 and 2.7.

Thales KeySecure server and CipherTrust Manager versions supported by various versions of IBM FlashSystem® A9000 and A9000R are shown in the following tables.
Important: FlashSystem A9000R microcode versions from 12.1.0 and up until 12.2.1.b do not support Thales KeySecure server.

Table 1. Thales KeySecure server support

IBM FlashSystem A9000 and A9000R version	Supported versions of Thales KeySecure server
12.0 - 12.0.2	8.3.2 or earlier

IBM FlashSystem A9000 and A9000R version	Supported versions of Thales KeySecure server
12.0.3	8.4 or earlier
12.2.1.b and later	8.6 - 8.9
12.3.x	8.6 - 8.11.x Note: Upgrading to 8.11 can only be done if 8.10 is first migrated from the global key-based certificate (chain) to a client certificate based authentication. For more information, see the IBM Redbooks® publication Data-at-rest Encryption for the IBM Spectrum® Accelerate Family.

Table 2. Thales CipherTrust Manager support

IBM FlashSystem A9000 and A9000R version	Supported versions of Thales CipherTrust Manager
12.3.x	• k170v 2.0.0 • k170v 2.2.x • k170v 2.3.x • k170v 2.6.x Note: For more information, see Configuring key servers.

Note: Future versions of Thales KeySecure server and CipherTrust Manager may also be supported.
SKLM is also used to locally create, distribute, backup, and manage the life cycle of keys and certificates. For the detailed information about the implementation and management of digital certificates, refer to the *Managing Digital Certificates across the Enterprise* Redbooks publication.

Internal encryption key management scheme

The internal encryption key management scheme does not warrant the purchase, deployment, or management of a dedicated, independent key management system, because the encryption key is generated and stored within the storage system. In addition, with an encryption internal encryption key management scheme, keys are not affected by software upgrades, and remain available upon the failure of up to two grid controllers.

User roles and permissions

User roles allow specifying which roles are applied and the various applicable limits.

Note: None of these system-defined users have access to data.

Table 1. User roles and permissions

User role	Permissions and limits	Typical usage
Read-only	Read-only users can only list and view system information.	The system operator, typically, but not exclusively, is responsible for monitoring system status and reporting and logging all messages.

Application administrator	Only application administrators carry out the following tasks:
	• Creating snapshots of assigned volumes
	• Mapping their own snapshot to an assigned host
	• Deleting their own snapshot
		Application administrators typically manage applications that run on a particular server. Application managers can be defined as limited to specific volumes on the server. Typical application administrator functions:
		• Managing backup environments:
		  ◦ Creating a snapshot for backups
		  ◦ Mapping a snapshot to back up server
		  ◦ Deleting a snapshot after backup is complete
		  ◦ Updating a snapshot for new content within a volume
		• Managing software testing environment:
		  ◦ Creating an application instance
		  ◦ Testing the new application instance

Security administrator	Carries out encryption-related tasks and does not have to be a storage administrator.	All encryption-related operations.
Storage administrator	Has permission to all functions, except: • Maintenance of physical components or changing the status of physical components • Only the predefined administrator, named *admin*, can change the passwords of other users A predefined storage administrator account is provided with the system and cannot be deleted.	Storage administrators are responsible for all administration functions.
Domain administrator	Can administer a specific domain. For more information, see Multi-tenancy.	When using the multi-tenancy feature.

User role	Permissions and limits	Typical usage
Technician	The technician is limited to the following tasks: • Physical system maintenance • Phasing components in or out of service A predefined technician account is provided with the system and cannot be deleted.	Technicians maintain the physical components of the system. Only one predefined technician is specified per system. Technicians are IBM technical support team members.

Note:

1. All users can view the status of physical components; however, only technicians can modify the status of components.
2. User names are case-sensitive.
3. Passwords are case-sensitive.

- **User groups**
  A user group is a group of application administrators who share the same set of snapshot creation permissions. This enables a simple update of the permissions of all the users in the user group by a single command.
- **Predefined users**
  Two administrative users are preconfigured on the storage system at the factory. These users cannot be deleted.
- **User information**
  Configuring users requires defining the following options:

User groups

A user group is a group of application administrators who share the same set of snapshot creation permissions. This enables a simple update of the permissions of all the users in the user group by a single command.

The permissions are enforced by associating the user groups with hosts or clusters. User groups have the following characteristics:

- Only users who are defined as application administrators can be assigned to a group.
- A user can belong to only a single user group.
- A user group can contain up to eight users.
- If a user group is defined with `access_all=yes`, application administrators who are members of that group can manage all volumes on the system.

Storage administrators create the user groups and control the various permissions of the application administrators.
Hosts and clusters can be associated with only a single user group. When a user belongs to a user group that is associated with a host, it is possible to manage snapshots of the volumes mapped to that host. User and host associations have the following properties:

- User groups can be associated with both hosts and clusters. This enables limiting application administrator access to specific volumes.
- A host that is part of a cluster cannot also be associated with a user group.
- When a host is added to a cluster, the associations of that host are broken. Limitations on the management of volumes mapped to the host is controlled by the association of the cluster.
- When a host is removed from a cluster, the associations of that host become the associations of the cluster. This enables continued mapping of operations so that all scripts will continue to work.

Predefined users

Two administrative users are preconfigured on the storage system at the factory. These users cannot be deleted.

Storage administrator
: This user provides the highest level of customer access to the system.
Predefined user name: admin
Default password: adminadmin

Technician
: This user is only for IBM® service personnel.
Predefined user name: technician
Default password: Password is predefined and is used only by the IBM technicians.

Note: Predefined users are always authenticated by the storage system, even if LDAP authentication has been activated for them.

User information

Configuring users requires defining the following options:

Role
: Specifies the role category that each user has when operating the system. The role category is mandatory (see User roles and permissions).

Name
: Specifies the name of each user allowed to access the system.

Password

All user-definable passwords are case sensitive.
Passwords are mandatory, can be 6 to 12 characters long, use uppercase or lowercase letters as well as the following characters: ~!@#$%^&*()_+-={}|:;<>?,./\[] .

Email
: Email is used to notify specific users about events through e-mail messages. Specifying the user's email address is optional and must adhere to the standard email format.

Phone and area code
: Phone numbers are used to send SMS messages to notify specific users about events. Phone number and area code stings can contain up to 63 digits, including hyphens (-) and periods (.)

User authentication and access control

The storage system features role-based authentication either natively or by using LDAP-based authentication.

The system provides:

Role-based access control
: Built-in roles for access flexibility and a high level of security according to predefined roles and associated tasks.

Two methods of access authentication
: The following methods of user authentication are supported:

  Native authentication
  : This is the default mode for authentication of users and groups that are defined on the storage system. In this mode, users and groups are authenticated against a database on the system.

  LDAP
  : When enabled, the system authenticates the users against an LDAP repository.

Note: The administrator and technician roles are always authenticated by the storage system, regardless of the authentication mode.

- **Native authentication**
  Native authentication if the default mode for authenticating users and user groups.
- **LDAP authentication**
  Lightweight Directory Access Protocol (LDAP) support enables the authentication of storage system users through an LDAP server repository, or directory.
- **LDAP authentication logic**
  The LDAP authentication procedure consists of several key steps.

Native authentication

Native authentication if the default mode for authenticating users and user groups.

In this mode, users and groups are authenticated against a database on the system, based on the submitted username and password, which are compared to user credentials defined and stored on the storage system.

The authenticated user must be associated with a user role that specifies the system access rights.

LDAP authentication

Lightweight Directory Access Protocol (LDAP) support enables the authentication of storage system users through an LDAP server repository, or directory.

When LDAP authentication is enabled, the username and password of a user accessing the storage system (through the CLI or GUI) are used by the system to log in to a specified LDAP directory. Upon a successful login, the storage system retrieves the user's group membership data stored in the LDAP directory, and then uses that information to associate the user with a storage system administrative role.

The group membership data is stored in a customer defined, preconfigured attribute on the LDAP directory. This attribute contains string values which are associated with administrative roles. These values might be LDAP Group Names, but this is not required by the storage system. The values the attribute contains, and their association with storage system administrative roles, are also defined by the customer.

Supported domains

The following LDAP authentication directories are supported:

- Microsoft Active Directory
- SUN directory
- Open LDAP

LDAP multiple-domain implementation

In order to support multiple LDAP servers that span over different domains, and in order to use the memberOf property, the storage system allows for more than one role for the Storage Administrator and the Read-Only roles.

The predefined administrative IDs “admin” and “technician” are always authenticated by the storage system, whether or not LDAP authentication is enabled.

Responsibilities division between the LDAP directory and the storage system

LDAP and the storage system divide responsibilities and maintained objects.

Following are responsibilities and data maintained by the storage system and the LDAP directory:

LDAP directory

- Responsibilities: user authentication for system users, and assignment of a system-related user group.
- Maintains: users, username, password, designated system-related LDAP groups.

Storage system

- Responsibilities: determining the appropriate user role by mapping the LDAP user group to a storage system role, and enforcement of any user access to the storage system.
- Maintains: mapping of the LDAP user group to a storage system role.

LDAP authentication logic

The LDAP authentication procedure consists of several key steps.

1. The LDAP server and system parameters must be defined.
2. A storage system user must be defined on that LDAP server. The storage system uses this user when searching for authenticated users. This user is later on referred to as system's configured service account.
3. The LDAP user requires an attribute in which the values of the storage system user roles are stored.
4. Mapping between LDAP user attributes and storage system user roles must be defined.
5. LDAP authentication must be enabled on the storage system.

Once LDAP is configured and enabled, the predefined user is granted with login credentials authenticated by the LDAP server, rather than the storage system itself.

Figure 1. Login to a specified LDAP directory

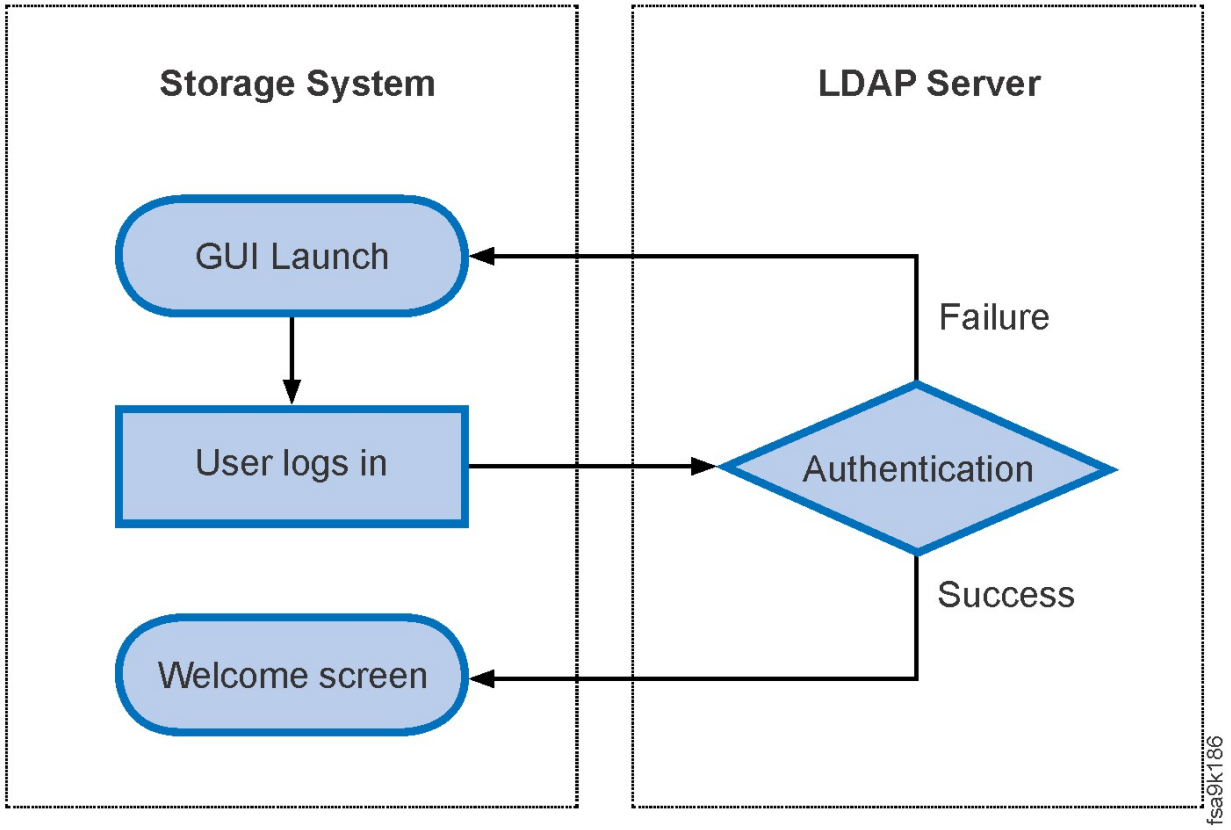

The storage administrator can test the LDAP configuration before its activation.

User validation

During the login, the system validates the LDAP user as follows:

Figure 2. The way the system validates users through issuing LDAP searches

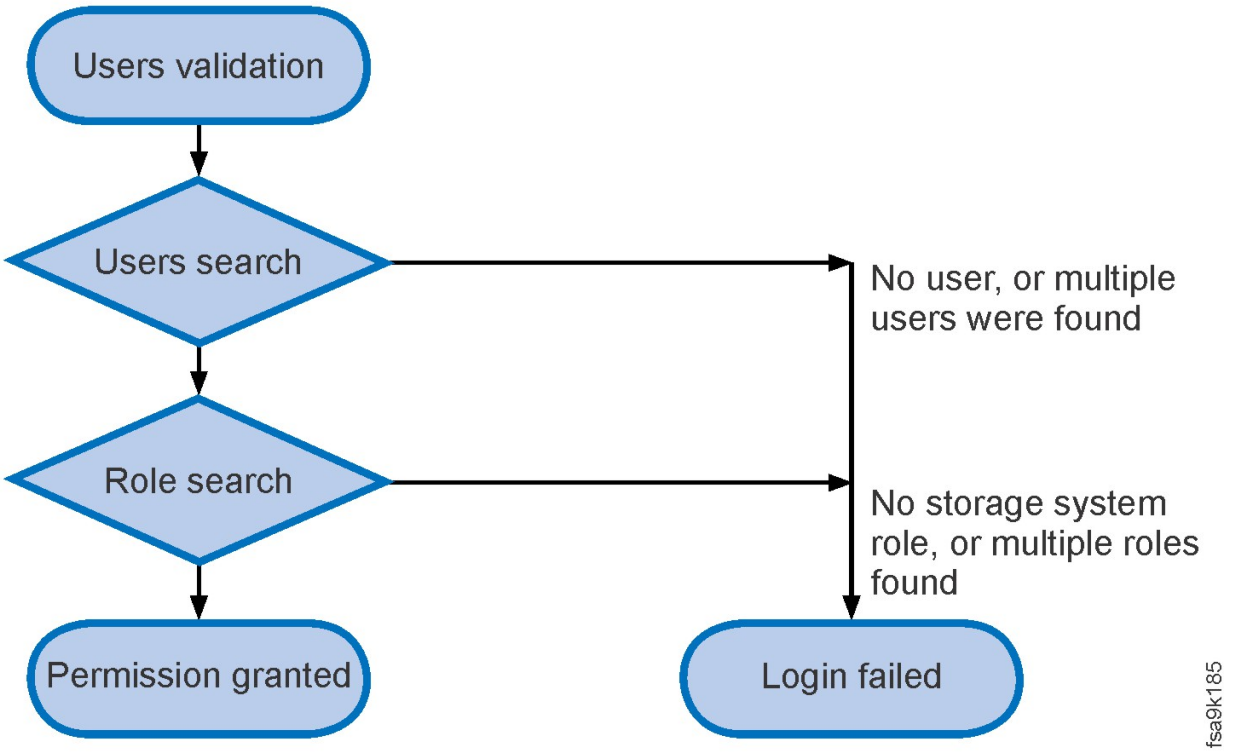

Issuing a user search

The system issues an LDAP search for the user's entered username. The request is submitted on behalf of the system's configured service account and the search is conducted for the LDAP server, base DN and reference attribute as specified in the storage system LDAP configuration.
The base DN specified in the storage system LDAP configuration serves as a reference starting point for the search – instructing LDAP to locate the value submitted (the username) in the attribute specified.

If a single user is found - issuing a storage system role search

The system issues a second search request, this time submitted on behalf of the user (with the user's credentials), and will search for storage system roles associated with the user, based on the storage system LDAP configuration settings.

- If a single storage system role is found - permission is granted

  The system inspects the rights associated with that role and grant login to the user. The user's permissions are in correspondence with the role associated by the storage system, based on the storage system LDAP configuration.
- If no storage system role is found for the user, or more than one role was found

  If the response by LDAP indicates that the user is either not associated with a storage system role (no user role name is found in the referenced LDAP attribute for the user), or is actually associated with more than a single role (multiple roles names are found) – login will fail and a corresponding message will be returned to the user.

If no such user was found, or more than one user were found

If LDAP returns no records (indicating no user with the username was found) or more than a single record (indicating that the username submitted is not unique), the login request fails and a corresponding message is returned to the user.

Multi-tenancy

The storage system allows allocating storage resources to several independent administrators, assuring that one administrator cannot access resources associated with another administrator.

Multi-tenancy extends the storage system approach to role-based access control. In addition to associating the user with predefined sets of operations and scope (the applications on which an operation is allowed), the storage system enables the user to freely determine what operations are allowed, and where they are allowed.

When multi-tenancy is enabled, storage resources are allocated to several independent administrators with the assurance that one administrator cannot view or access resources that are associated with another administrator.

This resource allocation is a partitioning of the system's resources to separate administrative *domains*.

A *domain* is a subset, or partition, of the system's resources. It is a named object to which users, storage pools, hosts, clusters, targets, and other entities can be associated. The domain restricts the resources that a user can manage to those associated with the domain.

A domain maintains the user relations that exist at the storage system level, as shown in the following figure.

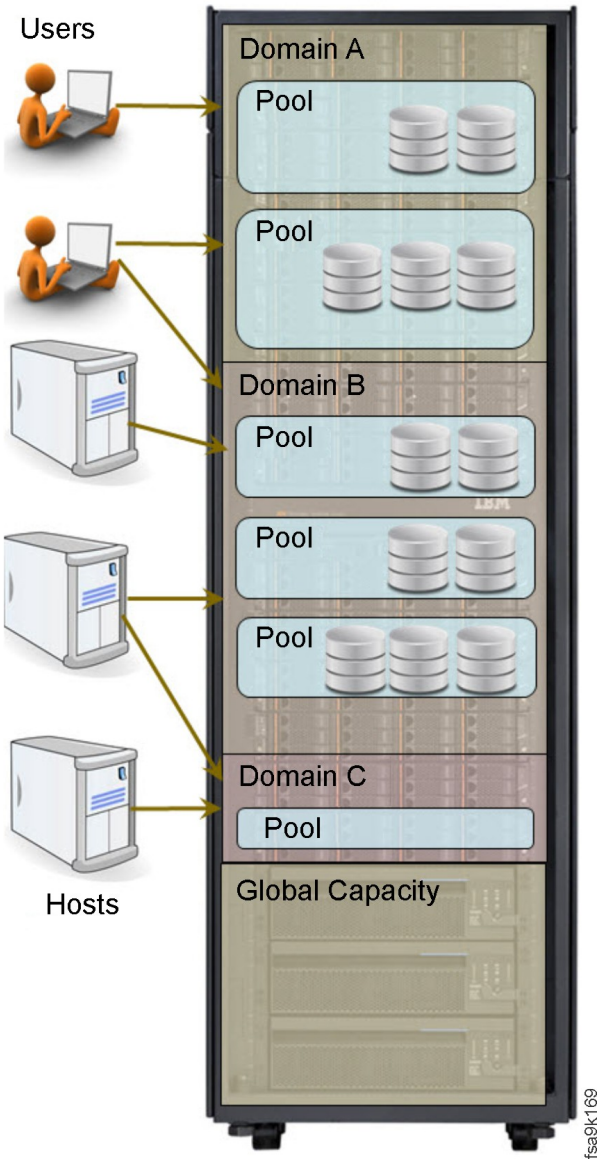

A *domain administrator* is the user who is associated with a domain and is able to manage it. The domain administrator is restricted to performing operations only on objects that are within the specific domain to which the administrator is associated.

The following access rights and restrictions apply to domain administrators:

- A user is created and assigned a role (for example: storage administrator, application administrator, read-only).
- When assigned to a domain, the user retains his given role, limited to the scope of the domain.
- Access to objects in a domain is restricted up to the point where the defined user role intersects the specified domain access.
- By default, domain administrators cannot access objects that are not associated with their domains.
- The domain is an isolated set of storage resources.
- The domain administrator has access only to the specified domains.
- The global administrator can manage domains and assign administrators to domains.
- Private objects are assigned to domains.
- The domain maintains its connectivity to global objects, such as: users, hosts, clusters, and targets. Hosts and clusters can serve several domains. However, hosts created by a domain administrator are assigned only to that domain.

Multi-tenancy offers the following benefits:

- **Partitioning** – The system resources are partitioned to separate domains. The domains are assigned to different tenants and each tenant administrator gets permissions for a specific, or several domains, to perform operations only within the boundaries of the associated domain(s).
- **Self-sufficiency** – The domain administrator has a full set of permissions needed for managing all of the domain resources.
- **Isolation** – There is no visibility between tenants. The domain administrator is not informed of resources outside the domain. These resources are not displayed on lists, nor are their relevant events or alerts displayed.
- **User-domain association** – A user can have a domain administrator role on more than one domain.
- **Users other than the domain administrator** – Storage, security, and application administrators, as well as read-only users, retain their right to perform the same operations that they have in a non-domain-based environment. These users can access the same objects under the same restrictions.
    - **Global administrator** – Not associated with any specific domain, and determines the operations that can be performed by the domain administrator in a domain. This is the only user that can create, edit, and delete domains, and associate resources to a domain.
      An *open* or *closed* policy can be defined so that a global administrator may, or may not, be able to see inside a domain.

      Intervention of a global domain administrator, that has permissions for the global resources of the system, is only needed for:

    - Initial creation of the domain and assigning a domain administrator.
    - Resolving hardware issues.
  - **Global user** – A user that is not associated with any domain has access rights to all of the entities that are not uniquely associated with a domain.

- **Working with multi-tenancy**
  This section provides a general description about working with multi-tenancy and its attributes.

Working with multi-tenancy

This section provides a general description about working with multi-tenancy and its attributes.

Domain administrator assignment and operations

The domain administrator assignment and permitted operations are as follows:

- Prior to its association with a domain, the future domain administrator (currently a system administrator) has access to all non-domain entities, and no access to domain-specific entities.
- When the storage administrator becomes a domain administrator, all access rights to non-domain entities are lost.
- The domain administrator can map volumes to hosts as long as both the volume and the host belong to the domain.
- The domain administrator can copy and move volumes across pools as long as the pools belong to domains administered by the domain administrator.
- Domain administrators can manage snapshots for all volumes in their domains.
- Domain administrators can manage consistency and snapshot groups for all pools in their domains. Moving consistency groups across pools is allowed as long as both source and destination pools are in the administrator's domains.
- Domain administrators can create and manage pools under the storage constraint associated with their domain.
- Although not configurable by the domain administrator, hardware list, and events - are available for view-only to the domain administrator within the scope of the domain.
- Commands that operate on objects not associated with a domain are not accessible by the domain administrator.

Domain attributes

The domain has the following attributes:

- **Capacity** – the domain is allocated with a capacity that is further allocated among its pools. The domain provides an additional container in the hierarchy of what was once *system-pool-volume*, and is now *system-domain-pool-volume*:
  - The unallocated capacity of the domain is reserved to the domain storage pools.
  - The sum of the physical capacity of the system's domains cannot exceed the storage system total physical capacity.
  - The sum of allocation limits of the system's domains cannot exceed the storage system maximum allocation limit.
- **Maximum number of volumes per domain** – the maximum number of volumes per system is divided among the domains in a way that one domain cannot consume all of the system resources at the expense of the other domains.
- **Maximum number of pools per domain** – the maximum number of pools per system is divided among the domains in a way that one domain cannot consume all of the system resources at the expense of the other domains.
- **Maximum number of mirrors per domain** – the maximum number of mirrors per system is divided among the domains.
- **Maximum number of consistency groups per domain** – the maximum number of consistency groups per system is divided among the domains.
- **Performance class** – the maximum aggregated bandwidth and IOPS is calculated for all volumes of the domain, rather than on a system level.
- The domain has a string that identifies it for LDAP authentication.

Mirroring a multi-tenancy environment

- The target, target connectivity and interval schedule are defined, edited and deleted by the storage administrator.
- The domain administrator can create, activate and change properties to a mirroring relation based on the previously defined target and target connectivity that are associated with the domain.
- The remote target does not have to belong to a domain.
- Whenever the remote target belongs to a domain, it checks that the remote target, pool and volume (if specified upon the mirror creation) all belong to the same domain.

Management and monitoring

The storage system can be monitored and fully controlled by using different management and automation tools.

The primary management tools for storage administrators are:

- **IBM® Hyper-Scale Manager** – Management server software connects to and controls one or more storage systems. Remote users can log into the server and use its advanced web-based user interface (UI) for managing and monitoring multiple storage systems in real time.
  The web UI can be accessed in a web browser pointing to the web address of the Hyper-Scale Manager server, together with the port number that is used to open the connection.

  For details, see IBM Hyper-Scale Manager documentation .

- **IBM XCLI Utility** – Provides a terminal-based command-line interface for issuing storage system management, monitoring, and maintenance commands from a client computer upon which the utility is installed.

The command-line interface is a comprehensive, text-based tool that is used to configure and monitor the system. Commands can be issued to configure, manage, or maintain the system, including commands to connect to hosts and applications.

For information on how to use IBM XCLI Utility, refer to the CLI Reference Guide.

Programmers can utilize the system's advanced application programming interfaces (APIs) for controlling and automating the system:

- Representational state transfer (REST) APIs
- CIM/SMI-S open APIs
- SNMP

Event reporting and handling

The health, configuration changes, and entire activity of the storage systems are constantly monitored by internal monitoring mechanisms that generate system events.

These events are accumulated by the system and can help the user in the following two ways:

- Users can view past events using various filters. This is useful for troubleshooting and problem isolation.
- Users can configure the system to send one or more notifications, which are triggered upon the occurrence of specific events. These notifications can be filtered according to the events, severity and code. Notifications can be sent through e-mail, SMS messages, or SNMP traps.

- **Event information**
  Each event has a predefined structure and is assigned a severity level.
- **Event notification rules**
  The storage system monitors the health, configuration changes, and activity of your storage systems and sends notifications of system events as they occur.
- **Event notification destinations**
  Event notifications can be sent to one or more destinations, meaning to a specific SMS cell number, e-mail address, or SNMP address, or to a destination group comprised of multiple destinations.
- **Event notification gateways**
  Some event notifications types require gateway definitions in order to enable the notification delivery.

Event information

Each event has a predefined structure and is assigned a severity level.

Events are created by various processes, including:

- Object creation or deletion, including volume, snapshot, map, host, and storage pool
- Physical component events
- Network events

Each event contains the following information:

- A system-wide unique numeric identifier
- A code that identifies the type of the event
- Creation timestamp
- Severity
- Related system objects and components, such as volumes, disks, and modules
- Textual description
- Alert flag, where an event is classified as alerting by the event notification rules.
- Cleared flag, where alerting events can be either uncleared or cleared. This is only relevant for alerting events.

Event information can be classified with one of the following severity levels:

Critical
: Requires immediate attention

Major
: Requires attention soon

Minor
: Requires attention within the normal business working hours

Warning
: Non-urgent attention is required to verify that there is no problem

Informational
: Normal working procedure event

The storage system provides the following variety of criteria for displaying a list of events:

- Before timestamp
- After timestamp
- Code
- Severity from a certain value and up
- Alerting events, meaning events that are sent repeatedly according to a snooze timer
- Uncleared alerts

The number of displayed filtered events can be restricted.

Event notification rules

The storage system monitors the health, configuration changes, and activity of your storage systems and sends notifications of system events as they occur.

Event notifications are sent according to the following rules:

Which events
: The severity, event code, or both, of the events for which notification is sent.

Where
: The destinations or destination groups to which notification is sent, such as cellular phone numbers (SMS), e-mail addresses, and SNMP addresses.

Notifications are sent according to the following rules:

Destination
: The destinations or destination groups to which a notification of an event is sent.

Filter
: A filter that specifies which events will trigger the sending of an event notification. Notification can be filtered by event code, minimum severity (from a certain severity and up), or both.

Alerting
: To ensure that an event was indeed received, an event notification can be sent repeatedly until it is cleared by a CLI command or GUI operation. Such events are called alerting events. Alerting events are events for which a snooze time period is defined in minutes. This means that an alerting event is resent repeatedly each snooze time interval until it is cleared. An alerting event is uncleared when it is first triggered, and can be cleared by the user. The cleared state does not imply that the problem has been solved. It only implies that the event has been noted by the relevant person who takes the responsibility for fixing the problem. There are two schemes for repeating the notifications until the event is clear: snooze and escalation.

Snooze
: Events that match this rule send repeated notifications to the same destinations at intervals specified by the snooze timer until they are cleared.

Escalation
: You can define an escalation rule and escalation timer, so that if events are not cleared by the time that the timer expires, notifications are sent to the predetermined destination. This enables the automatic sending of notifications to a wider distribution list if the event has not been cleared.

The following limitations apply to the configuration of alerting rules:

- Rules cannot escalate to nonalerting rules, meaning to rules without escalation, snooze, or both.
- Escalation time should not be defined as shorter than snooze time.
- Escalation rules must not create a loop (cycle escalation) by escalating to itself or to another rule that escalates to it.
- The configuration of alerting rules cannot be changed while there are still uncleared alerting events.

Event notification destinations

Event notifications can be sent to one or more destinations, meaning to a specific SMS cell number, e-mail address, or SNMP address, or to a destination group comprised of multiple destinations.

SMS destination

An SMS destination is defined by specifying a phone number. When defining a destination, the prefix and phone number should be separated because some SMS gateways require special handling of the prefix.

By default, all SMS gateways can be used. A specific SMS destination can be limited to be sent through only a subset of the SMS gateways.

E-mail destination

An e-mail destination is defined by an e-mail address. By default, all SMTP gateways are used. A specific destination can be limited to be sent through only a subset of the SMTP gateways.

SNMP managers

An SNMP manager destination is specified by the IP address of the SNMP manager that is available to receive SNMP messages.

Destination groups

A destination group is simply a list of destinations to which event notifications can be sent. A destination group can be comprised of SMS cell numbers, e-mail addresses, SNMP addresses, or any combination of the three. A destination group is useful when the same list of notifications is used for multiple rules.

Event notification gateways

Some event notifications types require gateway definitions in order to enable the notification delivery.

Email (SMTP) gateways

Several email gateways can be defined to enable notification of events by email. By default, the storage system attempts to send each email notification through the first available gateway according to the order that you specify. Subsequent gateways are only attempted if the first attempted gateway returns an error. A specific email destination can also be defined to use only specific gateways.

All event notifications sent by e-mail specify a sender whose address can be configured. This sender address must be a valid address for the following two reasons:

- Many SMTP gateways require a valid sender address or they will not forward the e-mail.
- The sender address is used as the destination for error messages generated by the SMTP gateways, such as an incorrect e-mail address or full email mailbox.

Email-to-SMS gateways

SMS messages can be sent to cell phones through one of a list of email-to-SMS gateways. One or more gateways can be defined for each SMS destination.

Each such email-to-SMS gateway can have its own SMTP server, use the global SMTP server list, or both.

When an event notification is sent, one of the SMS gateways is used according to the defined order. The first gateway is used, and subsequent gateways are only tried if the first attempted gateway returns an error.

Each SMS gateway has its own definitions of how to encode the SMS message in the email message.

Integration with ISV environments

The storage system can be fully integrated with different independent software vendor (ISV) platforms, APIs, and cloud environments, such as Microsoft Hyper-V, VMware vSphere, OpenStack, and more.

This integration can be implemented natively or by using IBM® cloud software solutions, which can facilitate and enhance this integration.

For more information about the available cloud storage solutions, see Platform and application integration.

- **Supporting VMware vStorage extended operations**
  Extended operations through the VMware vStorage APIs for Array Integration (VAAI) are natively supported, allowing the offload of these operations to the storage system.
- **Integration with Microsoft Azure Site Recovery**
  The Microsoft Azure Site Recovery (ASR) solution helps protect important applications by coordinating the replication and recovery of private clouds across remote sites.

Supporting VMware vStorage extended operations

Extended operations through the VMware vStorage APIs for Array Integration (VAAI) are natively supported, allowing the offload of these operations to the storage system.

The following extended operations are offloaded from the VMware ESXi server to the storage system:

Full copy (clone)
: Copies data from one logical unit (LUN based) to another without writing to the ESX server. Rather than issuing read and write requests from the host, the data copying operation is initiated on the storage system. This speeds up the virtual machine (VM) cloning operation and reduces the CPU load on the host.

Block zeroing
: Assigns zeroes to large storage areas without actually sending the zeros to the storage system. This speeds up the VM initiation operation, and reduces the I/O and CPU load on the host.

Hardware-assisted locking
: Locks a particular range of blocks in a shared logical unit, providing exclusive access to these blocks. Instead of using SCSI reservation that locks the entire logical unit, locking specific blocks is a more efficient alternative that greatly improves scalability in large server arrays. The locking is performed using Atomic Test & Set (ATS) commands.

The offloading of these extended operations from the ESXi server to the storage system reduces strain on the ESXi server, and saves a significant amount of time and computing resources when performing these operations in a full-scale production environment.

Integration with Microsoft Azure Site Recovery

The Microsoft Azure Site Recovery (ASR) solution helps protect important applications by coordinating the replication and recovery of private clouds across remote sites.

IBM FlashSystem® A9000 and A9000R can be integrated with Microsoft Azure Site Recovery, enabling customers using Microsoft System Center Virtual Machine Manager (SCVMM) to orchestrate and manage replication and disaster recovery. Support for Microsoft Azure Site Recovery is based on support for SMI-S (http://www.snia.org/ctp/conformingproviders/ibm.html#sftw7).

The SCVMM ASR integrates with storage solutions, such as the FlashSystem A9000 and A9000R CIM Agent, to provide site-to-site disaster recovery for Hyper-V environments by leveraging the SAN replication capabilities that are natively offered by the storage system. It orchestrates replication and failover for virtual

machines managed by SCVMM.

SCVMM ASR uses the remote mirroring feature (see Synchronous remote mirroring) through SMI-S to create and manage the replication groups.

Software upgrade

Non-disruptive code load (hot upgrade) upgrades the storage system software from a current version to a newer version without disrupting data access.

Important: Refer to the latest release notes for information on available upgrade paths for non-disruptive code load.
During upgrade process there is a point in time dubbed as the 'point-of-no-return', before which the process can still be aborted (either automatically by the system - or manually through a dedicated CLI). Once that point is crossed, the upgrade process is not reversible.

The upgrade process is run on all modules in parallel and is designed to be quick enough so that the applications' service on the hosts will not be impacted.

Following are notable characteristics of non-disruptive code load:

Duration of the upgrade process
: The overall process of downloading new code to storage system and moving to the new code is done non-disruptively.
The duration of the upgrade process may be affected by the system load during the upgrade. A very busy system may require a few trials until hot-upgrade is successfully completed.

Prerequisites and constraints

- As long as the overall hot upgrade is in progress (up to several minutes) no management operations are allowed (save for status querying), and no events are processed.
- All internal network paths must be active.

Effect on mirroring
: Mirrors are automatically deactivated for before the upgrade, and reactivated after it is over.

Effect on management operations
: During the upgrade process it is possible to query the system about the upgrade status, and the process can also be aborted manually before the 'point-of-no-return'. If a failure occurs before this point, the process is aborted automatically.

The upgrade is performed by IBM, and must be scheduled in advance. As of January 1st, 2019, the upgrade can be performed remotely. This optional process is called the Remote Code Load (RCL). It is recommended that a customer's Technical Advisor schedule the RCL, but customers can schedule it through the booking tool at https://a9000.ibmrcl.enterpriseappointments.com/v2/. For more information about the RCL process, see https://developer.ibm.com/storage/products/ibm-flashsystem/#rcl or speak to your Technical Advisor.

- **Preparing for software upgrade**
  The system code is upgraded without disconnecting active hosts or stopping I/O operations. However, some preparation is required.

Preparing for software upgrade

The system code is upgraded without disconnecting active hosts or stopping I/O operations. However, some preparation is required.

Attention: The upgrade must be performed only by an authorized IBM® service technician.
Before the code load (software upgrade), fulfill the following prerequisites by verifying that:

1. The native multipathing feature (provided by the operating system) is working on the connected hosts.
2. There are paths from the host to at least two different interface modules on the storage system.
3. There is no more than a single initiator in each zone (SAN Volume Controller attached to the system is an exception).
4. All hosts are connected using multipathing best practices. For more information, see Host network connectivity and multipathing configurations and Configuring host attachments.
5. In a case the storage system uses Fibre Channel (FC) connectivity for remote mirroring, the two systems should be connected to a SAN switch. Direct connection is not supported and is known to be problematic.
6. Hosts should be attached to the FC ports through an FC switch, and to the iSCSI ports through a Gigabit Ethernet switch. Direct attachment between hosts and to the storage system is not supported.

In addition, the following must be taken into consideration:

1. Co-existence with other multipathing software is not supported as GA (RPQ approval is required).
2. Connectivity to other storage servers from the same host is not supported as GA (RPQ approval is required).
3. Any remote mirroring connection is automatically suspended and resumed for a short while during the upgrade.
4. If a Microsoft Geo Cluster is used, check the cluster requirements.

Remote support and proactive support

To allow IBM® to provide support for the storage system, the proactive support and remote support options are available.

Note: For various preventive and diagnostics support actions relating to the storage system's continuous operation, IBM Support requires customer approval. Without customer approval, these support actions cannot be performed.

- **Proactive support** ("Call Home") – Allows proactive notifications regarding the storage system health and components to be sent to IBM Support at predefined intervals. Heartbeats and events are sent from the system to the IBM service center. The service center analyzes the information within the heartbeats and the events, correlates it with its vast database and can then trigger a component replacement prior to its potential failure.
  Upon detection of any hardware or software error code, both IBM Support and your predefined contact person are notified via e-mail, through a specified SMTP gateway. If IBM Support determines that the detected event requires service or further investigation, a new PMR is created and sent to appropriate IBM Support team. Proactive support minimizes the number of interaction cycles with IBM Support.
  Note: Proactive support is compliant with General Data Protection Regulation (GDPR) EU 2016/679. For more information, refer to http://ibm.biz/gdpr-storage-paper.
- If required, the customer email gateway can be configured to send call home information to IBM only via a secured channel. For more information, refer to the 'Encrypting Call Home and heartbeat notifications' section in the IBM FlashSystem A9000 and IBM FlashSystem A9000R Architecture and Implementation Redbooks publication.

- **Remote support** – Allows IBM Support to remotely and securely access your storage system when needed during a support call. This option requires IP communication between the storage system and the IBM Remote Support Center. If a storage system does not have direct access to the Internet (for example, due to a firewall), use the IBM Remote Support Proxy utility to enable the connection. Remote support minimizes the time it takes to diagnose and remedy storage system operational issues.

For more information, refer to the *Planning for remote support, on-site service, and maintenance* section of the *Deployment Guide*.

Deploying

Use the following information to prepare for the deployment of IBM FlashSystem® A9000R systems.

- **Deploying IBM HyperSwap Quorum Witness**
  The following sections describe how to install, manage, and troubleshoot IBM Spectrum Accelerate Family HyperSwap® Quorum Witness.

Roles and responsibilities

Both IBM® and the customer have roles and responsibilities that they must adhere to, in order to ensure proper workflow, timely successful installation, and properly configured Call Home and remote support, leading to a superior client experience.

Roles and responsibilities of the customer

- Review the product Deployment Guide.
- Enable and work with the Remote Support Center (RSC) remote support in performing remote data collection and support.
- Work with the IBM Planning Representative (IPR), Service Representative (SSR), Quality Practitioner (QPer), or other IBM personnel to fill out the Technical and Delivery Assessment (TDA) for an accurate and quicker initial installation.
- Provide and prepare a rack, adhering to the rack requirements, as specified in this guide.
- Provide adequate staffing/resources to support this solution.
- Provide sufficient bandwidth and host attachments to support this solution.
- Provide necessary Ethernet cabling.
- Provide all initial host Fibre Channel (FC) and iSCSI cabling.
- Provide proper power receptacles to match the requirements for the IBM FlashSystem® A9000R ordered.
- Provide proper thermal dissipation, airflow and cooling, and environmental requirements.
- Provide proper floor space and clearance.
- Provide access for the IBM service representative (SSR), including notebook computer or PC access.
- Provide access for movers and vehicles.
- Allow firewall access to Call Home servers.
- Setup IP host network.
- Setup SAN host networking.
- Install the Management Server from Fix Central for IBM Hyper-Scale Manager UI use.
- Perform logical configuration.
- Complete the host attachment plan.
- Prepare Fibre Channel (FC) connections.
- Prepare raised floor, if required.

CAUTION:

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

Roles and responsibilities of IBM Service Support Representatives (SSRs)

- Complete Distant Learning (DL) education and hands-on education course.
- Enroll in a hearing conservation program.
- Perform product installation.
- Configure Call Home and remote support.

- Install software upgrades.
- Install hardware Engineering Change Notices (ECA) also known as Field Bill of Materials (FBM).
- Conduct product relocation, at customer request.
- Perform break/fix repairs.
- Handle the return of failed parts that are under warranty or have a Certified Spare Parts Value.
- Keep customers informed of service activities.
- Arrange time with customer/TA to facilitate upgrades.
- Assist with break/fix support as requested by Remote Support Center, Top Gun, or PFE team member.
- Complete accurate Quality Service Activity Reporting (QSAR) reporting.

Note: Additional information can be found in your *Enterprise Class Support for Storage* document, provided by your IBM representative.

Overview

This guide defines deployment, planning, and preinstallation requirements for IBM FlashSystem® A9000R storage systems. It is important to ensure that you meet all requirements to help achieve a fast and reliable installation.

IBM FlashSystem A9000R is a grid-scale, all-flash storage platform designed to drive your business into the cognitive era.

FlashSystem A9000R provides consistent, efficient performance for dynamic data at scale. The FlashSystem A9000R storage system integrates the microsecond latency and high availability of IBM FlashCore® technology with grid architecture, comprehensive data reduction, and industry leading IBM® software.

Figure 1. IBM FlashSystem A9000R storage system

IBM FlashSystem A9000R is an excellent platform for industry leaders with rapidly growing cloud storage and mixed workload environments. IBM software-defined storage capabilities and IBM FlashCore technology combine to produce the extreme performance and scalability required in enterprise-class storage solutions.

The storage system utilizes IBM MicroLatency modules, providing density, low latency, high I/O, and high availability by leveraging IBM-enhanced MLC flash in model 415, and 3D triple-level cell (3D TLC) flash in models 425 and U25. IBM FlashSystem A9000R aggregates grid elements (each containing two grid controllers and one flash enclosure) within a 42U integrated rack solution. (For more information about grid elements, see FlashSystem A9000R grid elements).

Due to the combination of grid-scale architecture and flash storage media, the system delivers predictable high performance and ultra-low latency, even under heavy workloads with full data reduction enabled. As a result, the grid-scale architecture maintains this performance by automatically self-optimizing workloads across all storage resources without manual intervention. Secure multi-tenancy and quality of service (QoS) features help ensure that tenant service levels are not compromised within your complex environment.

IBM FlashSystem A9000R provides native implementation of IBM HyperSwap capability as well as Multi-site high availability and disaster recovery (HA/DR) capabilities. HyperSwap delivers active-active data access and transparent failover, per volume, across IBM FlashSystem® A9000 and IBM FlashSystem A9000R arrays and across data centers. Multi-site HA/DR supports high-availability with HyperSwap between two primary sites, and disaster recovery via simultaneous Asynchronous replication from the primary sites to a tertiary site.

IBM FlashSystem A9000R is ready for Kubernetes container environments and integrates with a wide variety of hypervisor and virtualization software, including IBM Bluemix, VMware, OpenStack, Linux, and Microsoft.

For more information regarding the IBM FlashSystem A9000R grid scale architecture, see Introduction > Architecture in the *IBM FlashSystem A9000R Product Overview* (SC27-8559).

- **Planning best practices and requirements**
  Use the planning information in this guide to place the FlashSystem A9000R system, plan power and environmental needs, plan for software and storage needs, and prepare for unique configurations that are based on how you plan to use the storage system.
- **Additional product information**
  This guide only covers deployment and planning information for the IBM FlashSystem A9000 storage system.

Planning best practices and requirements

Use the planning information in this guide to place the FlashSystem A9000R system, plan power and environmental needs, plan for software and storage needs, and prepare for unique configurations that are based on how you plan to use the storage system.

Good planning is essential for the successful setup and use of your IBM FlashSystem® A9000R. It ensures that you have everything you need and that you meet all specified system prerequisites. It also helps minimize errors and speeds up the installation process.

It is imperative that you work with the IBM sales team, IBM representative, and IBM Service Representative (SSR) to capture information needed to install and configure the storage system. This information is collected during a Technical and Delivery Assessment (TDA) or installation planning meeting. This information must be collected prior to the commencement of the installation, or delays may occur.
CAUTION:
⚠Customers must prepare their environments to handle the FlashSystem A9000R system based on this planning information with assistance from an IBM® representative. The final installation site within the computer room must be prepared *before* the equipment is delivered. If the site cannot be prepared before the delivery time, customers must make arrangements to have the professional movers return to finish the transportation later. Only professional movers can transport the equipment. The IBM service representative can minimally reposition the rack at the installation site, as needed to complete required service actions. Customers are also responsible for using professional movers in the case of equipment relocation or disposal.
If you cannot meet any of the installation requirements, notify your IBM service representative to help devise alternative solutions.

Additional product information

This guide only covers deployment and planning information for the IBM FlashSystem® A9000 storage system.

Additional product information

- For planning information for the IBM FlashSystem A9000 pod system, see *IBM FlashSystem A9000 Deployment Guide* , GC27-8564 on the IBM FlashSystem A9000 documentation .
- For information regarding the IBM Storage Utility Offering (SUO) (Model U25), see the IBM Storage Utility Offering at IBM Marketplace .

System specifications

This section provides information about the general properties, performance, physical features, storage capacity features, and host connectivity for the different IBM® FlashSystem™ A9000R storage system models.

General properties

Element	Model 415	Models 425 and U25
Grid controllers	Up to 12 active grid controllers, each containing: • Two Intel Xeon E5 v3 10-core 2.3 GHz processors • 384 GB DDR4 memory • Redundant battery backup units and power supply units	Up to 8 active grid controllers, each containing: • Two Intel Xeon E5 12-core 2.2 GHz processors • 384 GB DDR4 memory • Redundant battery backup units and power supply units
Software	IBM FlashSystem A9000 and A9000R software v12.0.x or later	IBM FlashSystem A9000 and A9000R software v12.2.x or later
Models and warranties	• 9835-415: 1-year warranty • 9837-415: 3-year enterprise-class warranty Warranties include onsite service, same day 24×7	• 9835-425: 1-year warranty • 9837-425: 3-year enterprise-class warranty • 9837-U25: 3-year enterprise-class warranty Warranties include onsite service, same day 24×7
Data reduction and efficiency	• Pattern removal • Global, inline deduplication • Inline compression • Space-efficient snapshots • Thin provisioning	
Encryption	Hardware-based AES-XTS 256-bit with centralized key management	
Backplane interconnect	InfiniBand	
Client operating system support	For a current list of platforms supported, please visit the IBM® System Storage® Interoperation Center (SSIC)	

Storage capacity features

Model 415

Element	Details									
Flash storage type	IBM-enhanced MLC									
	900 TB configuration					**1.8 PB configuration**				
Effective capacity[1] (TB)	300	450	600	750	900	600	900	1200	1500	1800
Maximum capacity[2] (TB)	1400	2000	2600	3000	3000	1400	2000	2600	3000	3000
Raw capacity (TB)	105.6	158.4	211.2	264.0	316.8	211.2	316.8	422.4	528	633.6
Flash enclosure type	Flash enclosure-150					Flash enclosure-300				

Element	Details									
Grid controllers	4	6	8	10	12	4	6	8	10	12
Flash enclosures	2	3	4	5	6	2	3	4	5	6
IBM MicroLatency® modules per flash enclosure	12 × 2.9 TB					12 × 5.7 TB				
[1]Typical effective capacity is the available capacity after system overhead (including over-provisioning and RAID protection) and after the data reduction benefits of pattern removal, deduplication, and compression. This assumes data reduction of up to a multiple of 5.26 to 1. [2]Maximum capacity refers to the effective capacity provisioning limit.										

Models 425 and U25

Element	Details											
Flash storage type	IBM-enhanced 3D TLC											
	720 TB configuration				**1700 TB configuration**				**3600 TB configuration**			
Effective capacity[1] (TB)	180	360	540	720	425	850	1275	1700	900	1800	2700	3600
Maximum capacity[2] (TB)	1200	2400	3600	4800	1200	2400	3600	4800	1200	2400	3600	4800
Physical capacity[3] (TB)	36	72	108	144	85	170	255	340	180	360	540	720
Raw capacity (TB)	55.3	110.6	166.1	221.2	129	258	387	516	258	516	774	1032
Grid controllers	3	4	6	8	3	4	6	8	3	4	6	8
Flash enclosures	1	2	3	4	1	2	3	4	1	2	3	4
IBM MicroLatency modules per flash enclosure	12 × 3.6 TB				12 × 8.5 TB				12 × 18 TB			
[1]Typical effective capacity is the available capacity after system overhead (including over-provisioning and RAID protection) and after the data reduction benefits of pattern removal, deduplication, and compression. This assumes data reduction of up to a multiple of 5 to 1. [2]Maximum capacity refers to the effective capacity provisioning limit. [3]Physical capacity is the available capacity after system overhead, including flash media over-provisioning and RAID protection.												

Physical features (all models)

Physical feature	Model 415		Models 425 and U25
Rack dimensions (H × W × D)	201.5 cm (42U) × 66.44 cm × 129.7 cm (79.3 in. × 25.4 in. × 51.1 in.)		
Front clearance	120 cm (47.2 in.)		
Rear clearance	100 cm (39.4 in.)		
Weight	Minimum configuration:	616 kg (1358 lbs)	Prior to October 2018 (4 grid controllers and 2 flash enclosures): 616 kg (1358 lbs) From October 2018 (3 grid controllers and 1 flash enclosure): 550 kg (1212 lbs)
	Maximum configuration:	935 kg (2061 lbs)	774 kg (1706 lbs)
Input voltage	200–240 V AC, 50/60Hz (+/-10% tolerance) via 30A–63A Note: Dependent on input phase configuration type (single, delta, or wye)		
Power usage	Minimum configuration:	3.085 kW (typical); 4.516 kW (max)	Prior to October 2018 (4 grid controllers and 2 flash enclosures): 3.085 kW (typical); 4.516 kW (max) From October 2018 (3 grid controllers and 1 flash enclosure): 1.914 kW (typical); 2806 kW (max)
	Maximum configuration:	8.81 kW (typical); 13.91 kW (max)	5.996 kW (typical); 8.671 kW (max)
	See Power consumption.		

Operation environment (all models)

Environment	Details
Temperature range	10 - 35° C (50 - 95° F)
Maximum altitude	2134 m (7000 ft.)
Humidity	25 - 80% non-condensing

Host connectivity (all models)

Host connectivity feature		Details
Host system interfaces (per grid controller)	Storage systems with Fibre Channel (FC) capabilities)	4 × 16 Gb Fibre Channel + 2 × 10 Gb iSCSI
	Storage systems with iSCSI (Ethernet) capabilities only	4×10 Gb iSCSI
Host connectivity for mirroring (per system)	Recommended minimum link bandwidth value	50 Mbps
	Recommended maximum round trip latency value	250 ms

Host connectivity feature		Details
	Attaching IBM FlashSystem® A9000 and A9000R systems for mirroring	The connection between two FlashSystem A9000R systems or between a FlashSystem A9000R and a FlashSystem A9000 system has to pass through: • Ethernet LAN for iSCSI connections • SAN for FC connections
Note: In model 415, synchronous remote mirroring is supported by version 12.0.1 and later.		

Physical configuration options

Use these general guidelines for determining and ordering the feature codes that you need to customize your IBM FlashSystem® A9000R system.

Procedure

Note: Contact your IBM representative to help determine which configuration options are best for your needs.
To determine the required ordering information, answer the following questions:

1. Which model best fits your warranty requirements?
   See full warranty information detailed on the IBM FlashSystem A9000R documentation .
2. What are your capacity needs?
   See Storage capacity features for full capacity specifications and see Flash enclosure components and feature codes for flash enclosure information and feature codes.
3. What are your performance requirements?
   See Rack configurations to help determine how many grid elements you require.
4. What type of host connectivity do you need?
   See Host connectivity (all models) for full host connectivity specifications and see Grid controller components and feature codes for grid controller information and feature codes.
5. Do you require any of the following:
   - Water cooling?
     See Rear-door heat exchanger.
   - Weight or height reduced shipping?
     See Weight reduced shipping option and Height reduced shipping option.
   - A radio frequency identification device (RFID) tag?
     See Radio frequency identification device option.

- **Components and interconnect**
  This section lists the components and interconnection options that are supplied with each IBM FlashSystem A9000R integrated rack.
- **Rack configurations**
  IBM FlashSystem A9000R configurations are ordered according to your performance and capacity needs. Valid configurations vary according to model. Use the information in this section to help determine the best configuration for your needs.
- **Flash enclosure components and feature codes**
  The flash enclosures are used for IBM FlashSystem A9000R storage functions.
- **Grid controller components and feature codes**
  Each grid controller contains two hard disk drives (HDDs) and two solid state drives (SSDs) for IBM FlashSystem A9000R performance functions.
- **Rear-door heat exchanger**
  The rear-door heat exchanger (feature code AFR1) is an optional water-cooled device that is mounted on the rear of an IBM FlashSystem A9000R system. It cools the air that is heated and exhausted by devices inside the rack.
- **Height reduced shipping option**
  This information describes the Height Reduction shipping option for IBM FlashSystem A9000R, feature code AFR2.
- **Weight reduced shipping option**
  This information describes the Weight Reduction shipping option for the storage system, feature code AFR3.
- **Radio frequency identification device option**
  IBM offers an optional radio frequency identification device (RFID) for the storage system, feature code AFR5.

Components and interconnect

This section lists the components and interconnection options that are supplied with each IBM FlashSystem® A9000R integrated rack.

Table 1 breaks down the various components and subcomponents of the IBM FlashSystem A9000R.

Table 1. Components and interconnection options in IBM FlashSystem A9000R

Component	Model 415	Models 425 and U25	Subcomponents
IBM® T42 Enterprise rack			

Component	Model 415	Models 425 and U25	Subcomponents

Grid controllers	4 - 12, depending on ordered configuration	3 - 8, depending on ordered configuration	Each grid controller includes:
			• Two hot-swappable power supply units (PSUs).
			• Two internal battery modules.
			• Two data reduction hardware acceleration cards.
			• Two or three dual-ported host adapters, configuration according to customer request. Options include:
			  ◦ Two 16 Gb Fibre Channel dual ports (total four ports) and a 10 Gb Ethernet (iSCSI) dual port (total two ports) adapter *or*
			  ◦ Two 10 Gb Ethernet (iSCSI) dual port (total four ports) adapters

Flash enclosures	2 - 6, depending on ordered configuration	1 - 4, depending on ordered configuration	• **Model 415:** 12 hot-swap 2.9, or 5.7 TB IBM MicroLatency® modules **Models 425 and U25:** 12 hot-swap 3.6, 8.5, or 18 TB IBM MicroLatency modules • Two internal battery modules
Two InfiniBand switches			
Two power distribution units (PDUs)			
Utility patch panel			• Three management ports • One or two VPN ports (configuration dependent)
Internal cabling			

Note: The maintenance module previously provided with the 983x-415 is obsolete and no longer provided with this rack configuration.

Rack configurations

IBM FlashSystem® A9000R configurations are ordered according to your performance and capacity needs. Valid configurations vary according to model. Use the information in this section to help determine the best configuration for your needs.

Performance and capacity vary according to rack configuration:

- The more grid controllers in the system:
  - Increases the amount of compute resources and provides better performance.
  - Increases the number of network ports and provides additional bandwidth.

  For more information on the different network options for the grid controllers, see Grid controller components and feature codes.
- The more flash enclosures in the system increases capacity in the system.
  For more information on the different capacity options for the flash enclosures, see Flash enclosure components and feature codes.
- The IBM® MicroLatency® module raw capacity size ordered.
  Select the flash enclosure feature code that best suits your needs. Larger MicroLatency modules provide higher capacity at constant performance.

  For more information on the different capacity options for the flash enclosures, see Flash enclosure components and feature codes.

For various configuration feature codes see:

- Flash enclosure components and feature codes
- Grid controller components and feature codes
- Main power cables specifications

For more information regarding capacity, processors, memory, and connectivity, see System specifications.

Common components

All FlashSystem A9000R configurations include the following common components:

- Two InfiniBand switches
- Two power distribution units (PDUs)
- One utility patch panel

These components occupy a total of 5U in the rack.
Important: All space in the rack is to be left solely for the use of the IBM FlashSystem A9000R system. Do not use this rack for any other purpose.

- **FlashSystem A9000R grid elements**
  A *grid element* is the term used to describe a combination of one flash enclosure and two grid controllers within FlashSystem A9000R machines.
- **Model 415**
  Model 415 rack configurations have a range between a minimum of two grid elements to a maximum of six grid elements.
- **Models 425 and U25**
  Model 425 and U25 rack configurations have a range between a minimum configuration of three grid controllers and one flash enclosure (one grid element with one additional grid controller) to a maximum of four grid elements.

- **Conducting system scale-out (MES)**
  System scale-out (MES) is a performance and capacity expansion option for your system rack, if your system is not already fully populated. Performing system scale-out is a non-disruptive process.

FlashSystem A9000R grid elements

A *grid element* is the term used to describe a combination of one flash enclosure and two grid controllers within FlashSystem A9000R machines.

With the exception of model 425 and U25 minimum configurations starting October 2018 (see Models 425 and U25), all FlashSystem A9000R grid element valid configurations contain a number of flash enclosures and double that number of grid controllers.

For example, a three grid element configuration contains three flash enclosures and six grid controllers.
Note:

- A single grid element is **not** a functionally separate subsystem. In any single system each of the system grid controllers access each of the system flash enclosures.
- The size of each individual flash enclosure and grid controller is 2U. Hence, each grid element occupies 6U in the rack.

For more information on the various types of rack configurations, see Model 415 and Models 425 and U25.

Model 415

Model 415 rack configurations have a range between a minimum of two grid elements to a maximum of six grid elements.

A *minimum rack* configuration is comprised of two grid elements and common components. *Partially populated* rack configurations contain three, four, or five grid elements and common components. The maximum configuration (*full rack*) is comprised of six grid elements and common components.

Figure 1 shows a model 415 minimum rack configuration. Figure 2 shows a model 415 full rack configuration.

Figure 1. Model 415 minimum rack configuration

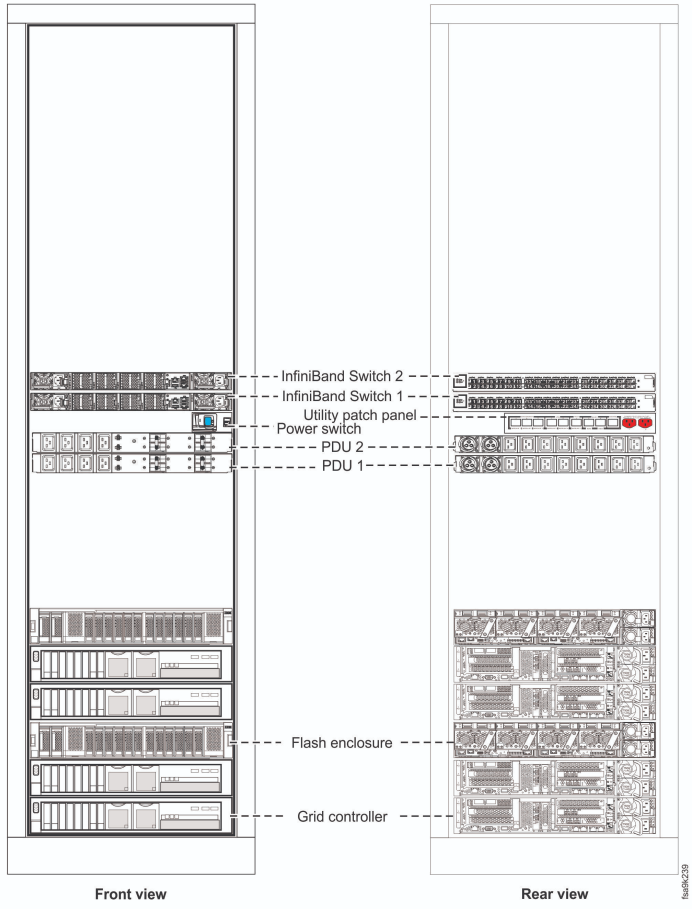

Figure 2. Model 415 full rack configuration

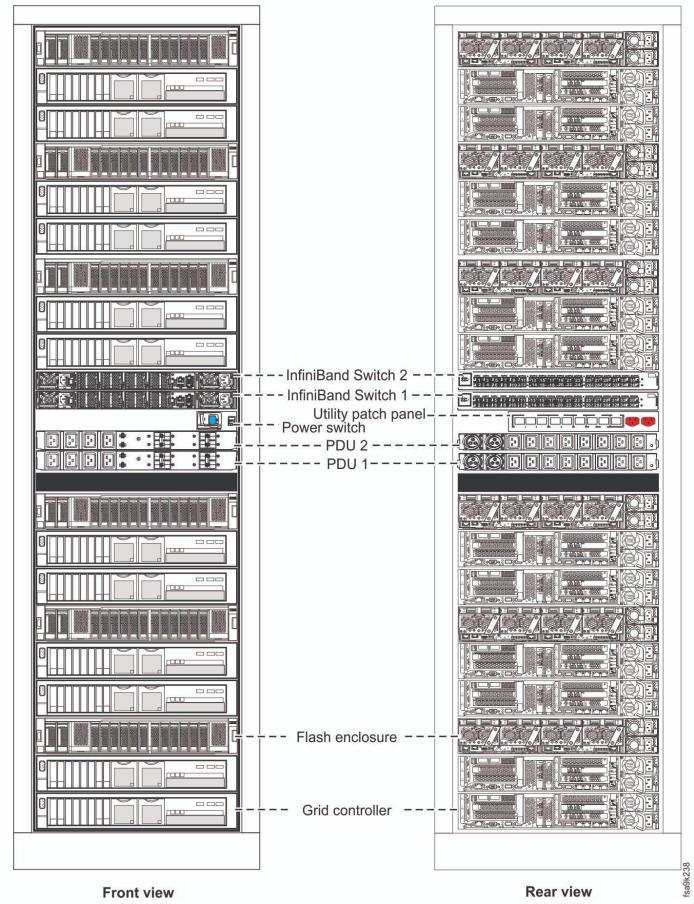

All flash enclosures contain 12 IBM® MicroLatency® modules. Depending on the flash enclosure ordered, MicroLatency modules are available with 2.9 TB or 5.7 TB raw capacity. For more information about flash enclosures components and feature codes, see Flash enclosure components and feature codes.

Models 425 and U25

Model 425 and U25 rack configurations have a range between a minimum configuration of three grid controllers and one flash enclosure (one grid element with one additional grid controller) to a maximum of four grid elements.

A *minimum rack* configuration is comprised of three grid controllers and one flash enclosure (one grid element and one additional grid controller) and common components. *Partially populated* rack configurations contain two or three grid elements and common components. The maximum configuration (*full rack*) is comprised of four grid elements and common components.

Note: Some regions only support a maximum of three grid elements. For more information, contact your IBM representative.
Minimum rack configurations contain one additional grid controller than the standard single grid element. This configuration ensures that a single grid controller failure does not result in a system shutdown.
Figure 1 shows a model 425 and U25 minimum rack configuration. Figure 2 shows a model 425 and U25 full rack configuration.

Figure 1. Models 425 and U25 minimum rack configuration

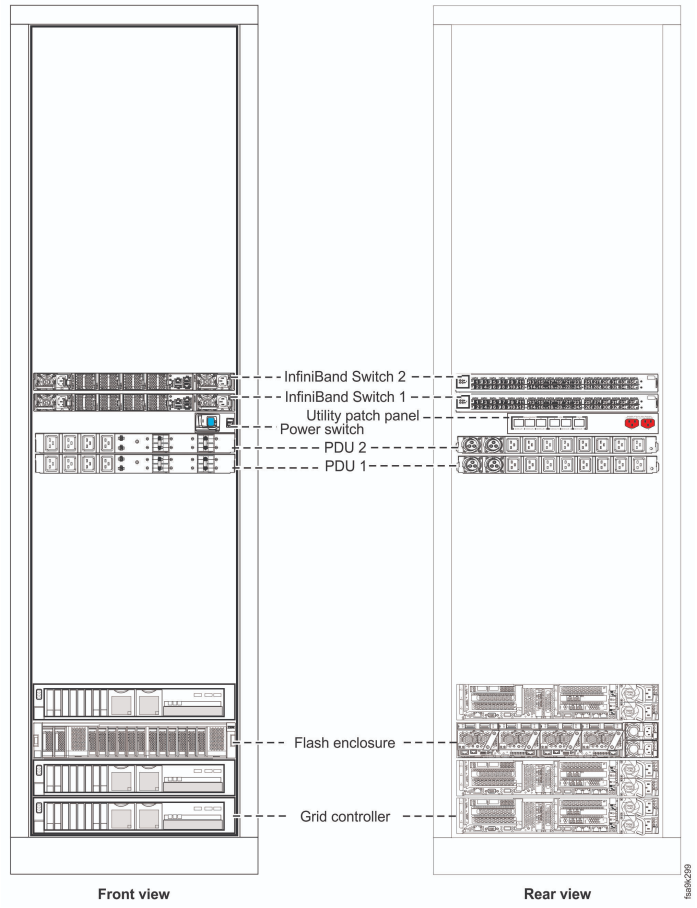

Figure 2. Models 425 and U25 full rack configuration

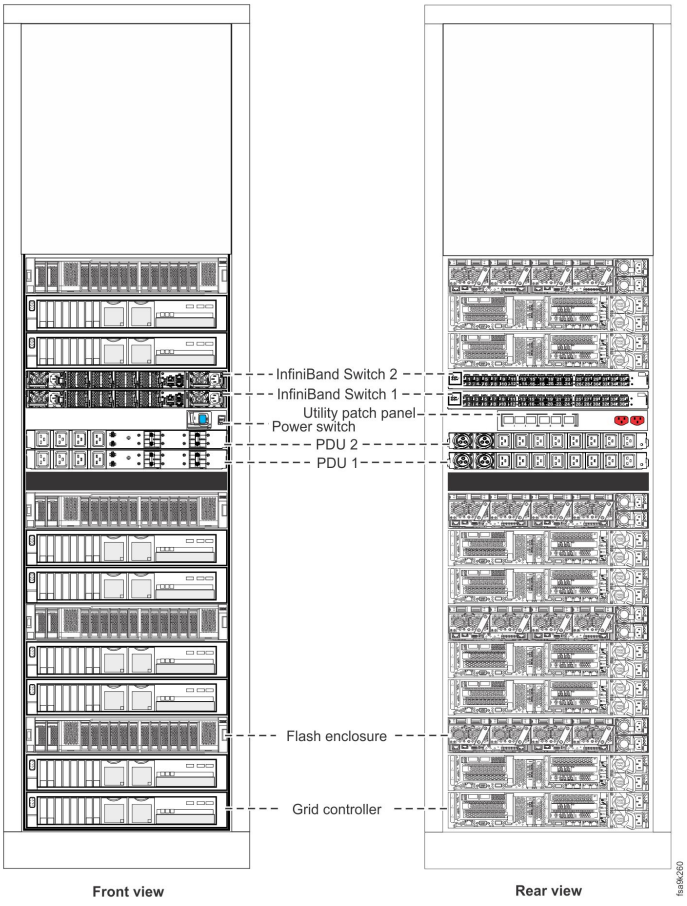

All flash enclosures contain 12 IBM® MicroLatency® modules. Depending on the flash enclosure ordered, MicroLatency modules are available with 3.6 TB, 8.5 TB, or 18 TB raw capacity. For more information about flash enclosures components and feature codes, see Flash enclosure components and feature codes.

Conducting system scale-out (MES)

System scale-out (MES) is a performance and capacity expansion option for your system rack, if your system is not already fully populated. Performing system scale-out is a non-disruptive process.

Table 1 shows which system versions of each FlashSystem A9000R model and configuration support system scale-out.

Table 1. System scale-out availability

Model	Configuration	Earliest system version allowing system scale-out
415	All configurations	12.0.3
425 and U25	2 grid elements and up (minimum of 4 grid controllers and 2 flash enclosures)	12.2.1
425 and U25	Minimum configuration (3 grid controllers and 1 flash enclosure)	12.3.1

When conducting system scale-out, additional grid controllers and flash enclosures are added to the system according to model configuration (see FlashSystem A9000R grid elements).

In order for the installation site to support any future capacity upgrade procedures, your site should adhere to all requirements listed in this document for a full rack. These requirements include:

- Floor-load requirements
- Power requirements

For more information regarding system scale-out ordering and requirements, contact your IBM representative.

Flash enclosure components and feature codes

The flash enclosures are used for IBM FlashSystem® A9000R storage functions.

Figure 1 shows the front view of a flash enclosure.

Figure 1. Front of a flash enclosure

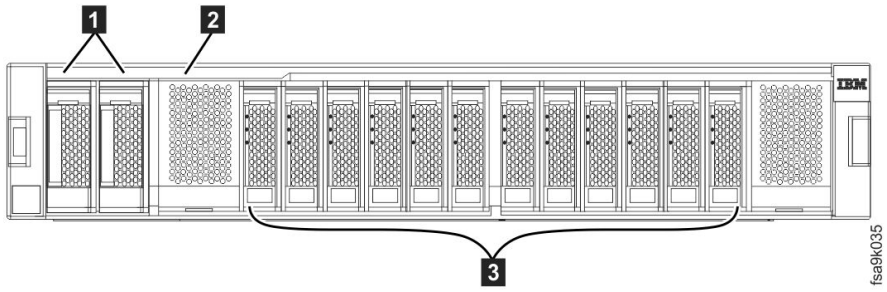

- 1 Battery modules
- 2 LED indicator panel
- 3 IBM® MicroLatency® modules

Figure 2 shows the rear view of a flash enclosure.

Figure 2. Rear of a flash enclosure

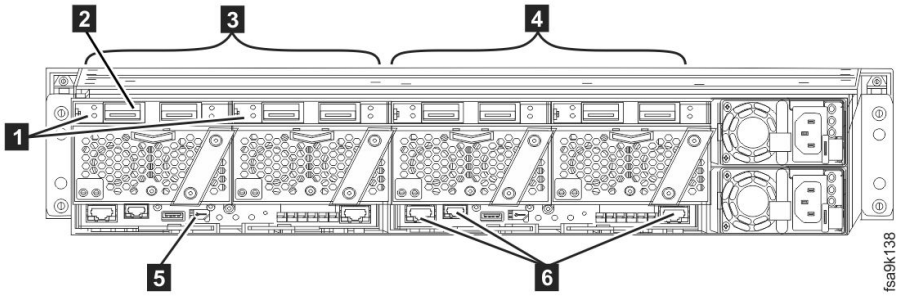

- 1 InfiniBand Adapters (two per canister)
- 2 Internal interconnect Infiniband ports
- 3 Canister 1
- 4 Canister 2
- 5 Maintenance ports (serial) (IBM technician access)
- 6 Management ports (Ethernet) (IBM technician access)

Model 415

Each rack unit contains 2 - 6 flash enclosures, according to customer specifications. Every flash enclosure contains 12 hot-swap 2.9 TB or 5.7 TB IBM MicroLatency modules.

Figure 1 and Figure 2 illustrate the front and rear of the flash enclosure.

Table 1 lists the feature codes for model 415.

Table 1. Feature codes for flash enclosures – model 415

Description	Feature code
Flash enclosure with 12 x 2.9 TB IBM MicroLatency modules	AFE2
Flash enclosure with 12 x 5.7 TB IBM MicroLatency modules	AFE3

Models 425 and U25

Each rack unit contains 1 - 4 flash enclosures, according to customer specifications. Every flash enclosure contains 12 hot-swap 3.6 TB, 8.5 TB, or 18 TB IBM MicroLatency modules.

Figure 1 and Figure 2 illustrate the front and rear of the flash enclosure.

Table 2 lists the feature codes for models 425 and U25.

Table 2. Feature codes for flash enclosures – models 425 and U25

Description	Feature code
Flash enclosure with 12 x 3.6 TB IBM MicroLatency modules	AFE4
Flash enclosure with 12 x 8.5 TB IBM MicroLatency modules	AFE5
Flash enclosure with 12 x 18 TB IBM MicroLatency modules	AFE6

Grid controller components and feature codes

Each grid controller contains two hard disk drives (HDDs) and two solid state drives (SSDs) for IBM FlashSystem® A9000R performance functions.

Figure 1 illustrates the front of the grid controller. Figure 2 and Figure 3 illustrate the two rear options for the grid controller.

Figure 1. Front of a grid controller

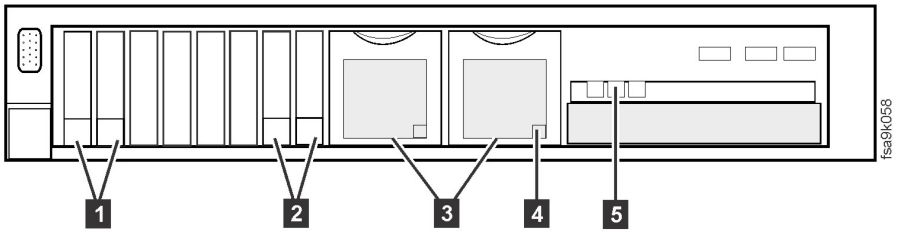

- 1 HDDs (two per grid controller)
- 2 SSDs (two per grid controller)
- 3 Battery modules (two per grid controller)
- 4 Battery module Fault LED (one per unit)
- 5 General battery module LEDs

Attention: **Do not** use the power button / LED (not shown in Figure 1) to turn off the machine. This can lead to a customer impact event (CIE).
For full information on how to properly shutdown a system, see the *IBM FlashSystem A9000R Command-Line Interface (CLI) Reference Guide* (SC27-8711).

Figure 2. Rear of a grid controller with FC configuration

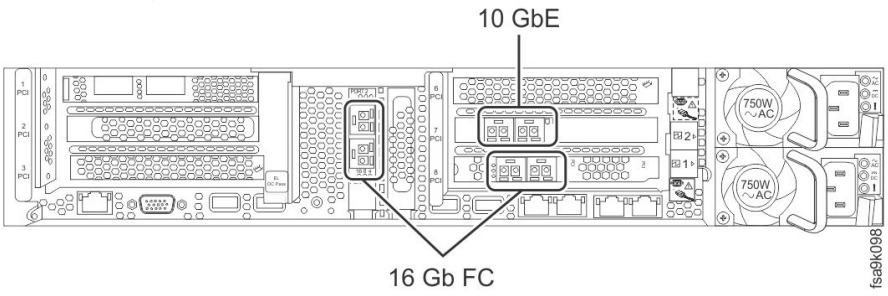

Figure 3. Rear of a grid controller with 10 Gb Ethernet configuration

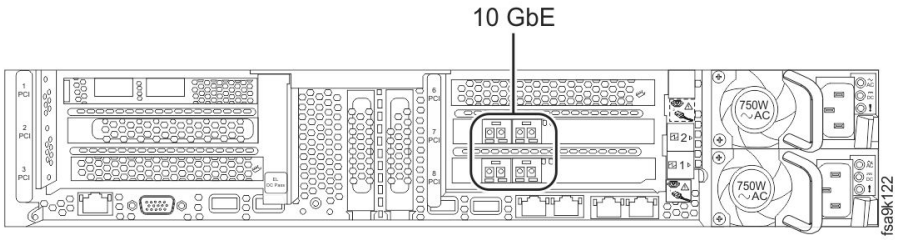

Model 415

Each rack unit contains 4 - 12 grid controllers, according to customer specifications.

Table 1 lists the feature codes for model 415.

Table 1. Feature codes for grid controllers – model 415

Description	Feature code
Grid controller with 4 x 16 Gb Fibre Channel (FC) + 2 x 10 GbE iSCSI	5001
Grid controller with 4 x 10 GbE iSCSI	5002

Models 425 and U25

Each rack unit contains 3 - 8 grid controllers, according to customer specifications.
Table 2 lists the feature codes for models 425 and U25.

Table 2. Feature codes for grid controllers – models 425 and U25

Description	Feature code
Grid controller with 4 x 16 Gb Fibre Channel (FC) + 2 x 10 GbE iSCSI	5003
Grid controller with 4 x 10 GbE iSCSI	5004
Grid controller with 4 x 16 Gb Fibre Channel NVMe/FC + 2 x 10 GbE iSCSI	5005

Rear-door heat exchanger

The rear-door heat exchanger (feature code AFR1) is an optional water-cooled device that is mounted on the rear of an IBM FlashSystem® A9000R system. It cools the air that is heated and exhausted by devices inside the rack.

A supply hose delivers chilled, conditioned water to the heat exchanger. A return hose delivers warmed water back to the water pump or chiller (referred to as the secondary cooling loop).

The primary cooling loop supplies the building chilled water to secondary cooling loops and air conditioning units.
Note: The hoses for the secondary cooling loop are not included with this option.
The rack on which you install the heat exchanger can be on a raised floor or a non-raised floor.

The rear-door heat exchanger option consists of the following components that are shown in Figure 1.

- Door assembly
- Hinge kit
- Air-purge tool

Figure 1. Rear-door heat exchanger option kit

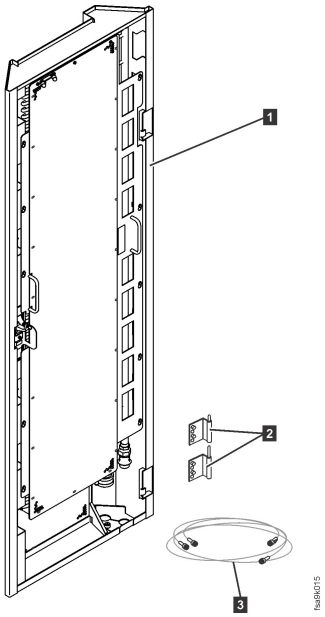

- 1 Door assembly
- 2 Hinge kit
- 3 Air-purging tool

The following table lists the rear-door heat exchanger feature code for the FlashSystem A9000R system.

Table 1. Rear-door heat exchanger features

Feature description	Feature code
Rear-door heat exchanger	AFR1

See Preparing for the rear-door heat exchanger for information about the requirements for preparing the installation site before the rear-door heat exchanger feature can be installed.

Height reduced shipping option

This information describes the Height Reduction shipping option for IBM FlashSystem® A9000R, feature code AFR2.

If your site does not meet the delivery clearances that are shown in Rack dimensions and service clearance requirements, the height reduced shipping option can be ordered to reduce the height of the rack by 30 cm (11.8 in.).

After the rack is delivered, the IBM® service representative removes the rack top cover so that the rack can be moved to the final location. Only professional movers can transport the equipment.

After the rack is in its final location, the IBM service representative must return to complete the installation, including reinstalling the rack top cover.

A fully configured FlashSystem A9000R system, including packaging, weighs approximately 935 kg (2061 lbs) (model 415) or 774 kg (1706 lbs) (model 425) with dimensions of 66 cm × 118 cm (26 in. × 46.5 in.).
Attention: The storage system must not be tilted more than 10 degrees.
Note: This option greatly increases the system installation time. Onsite coordination is needed for both the IBM service representative and the professional movers.

Weight reduced shipping option

This information describes the Weight Reduction shipping option for the storage system, feature code AFR3.

IBM® offers weight reduced shipping for all IBM FlashSystem® A9000R system configurations. This optional feature provides that the weight of the rack is only approximately 423 kg (932 lb) for traversal at the delivery site.

This option is ordered for installations where receiving an assembled storage unit, or ordering feature code AFR2 would be impractical, due to greater weight limitations.

The unit is delivered fully tested but partially disassembled into several easily transported subassemblies. This feature allows racks to be safely transported on lower weight capacity elevators.

At the installation site the system is unpacked by IBM technicians, and the front door, rear door and side covers are temporarily removed so that the rack assembly can then be carried up stairs, hoisted via crane through windows, tipped to fit through low doorways, and rolled through low doorways in the customer's facility.

At the final destination, all devices that were shipped separately can be installed.

The following are the approximate weights, including packaging, of a fully configured FlashSystem A9000R system:

- Model 415: 935 kg (2061 lbs)
- Models 425 and U25: 774 kg (1706 lbs)

Attention: The storage system must not be tilted more than 10 degrees.
Note: This option greatly increases the system installation time. Onsite coordination is needed for both the IBM service representative and the professional movers.

Radio frequency identification device option

IBM® offers an optional radio frequency identification device (RFID) for the storage system, feature code AFR5.

If you use frequency identification device (RFID) technology to track equipment in your data centers, you can order the RFID option to attach an RFID tag to system racks.

This RFID is designed to meet the performance and numbering specification as outlined by the radio frequency identification specifications. For more information about RFID, RFID numbering specifications, and general RFID best practices, see the *FSTC RFID IT Assets – Phase 2 Business Case and Best Practices* report, published by the Financial Services Technology Consortium (FSTC). This report can be found on the Publications page of the IBM FlashSystem® A9000R documentation .

Important: This option is applicable only in environments that can use the correct RFID reading technology. Before you order this option, review the RFID capabilities with your IBM service representative.

When this option is ordered, IBM attaches one RFID tag per rack. Order one RFID option for each FlashSystem A9000R that you want to track. This option does not tag individual components.
Important: If the tag must be replaced for an IBM FlashSystem A9000R system, ensure that you update the asset-management database with the new RFID number for that FlashSystem A9000R.

Physical installation site requirements

The location where you plan to install the storage system must meet all physical requirements.

Procedure

Note: Contact your IBM representative to help determine which configuration options are best for your needs.
Prepare the site in advance so that professional movers or riggers can transport the equipment to the final site within the computer room. If the site cannot be prepared before the delivery time, you must make arrangements to have the professional movers return to finish the transportation later.
Attention: Only professional movers should transport the equipment.
An IBM® service representative installs the storage system. The IBM service representative can only minimally reposition the rack within the room, as needed to complete required service actions.

Professional movers or riggers are required to transport the FlashSystem A9000R rack as close to the installation site as possible because of its weight.
Note: Professional movers or riggers are also required to relocate or dispose of the FlashSystem A9000R system.
To determine the required ordering information, answer the following questions regarding the installation site:

1. What type of power input do you have?
   See Input voltages and frequencies.
2. Does the installation site meet the physical site requirements for the FlashSystem A9000R and features that you plan to order?
   For example:
   - Can you space the racks to allow for sufficient floor strength? See Floor and space requirements.
   - Is there adequate cooling capacity to handle the new equipment? See Air circulation and cooling.
   - Is sufficient power available? See Power consumption.
   - Do you need 30A or 60A power cables? See Main power cables specifications.
   - If water cooling is required, see Planning for the rear-door heat exchanger.

- **Floor and space requirements**
  Ensure that the location of the FlashSystem A9000R system meets floor and space requirements.
- **Power requirements**
  Ensure that your operating environment meets the AC-power and voltage requirements for IBM FlashSystem A9000R systems.
- **Environmental requirements**
  Ensure that the installation site meets all operating environment requirements for the IBM FlashSystem A9000R system.
- **Planning for the rear-door heat exchanger**
  If the optional rear-door heat exchanger (feature code AFR1) was ordered, you must prepare the installation site before an IBM service representative can install the rear-door heat exchanger on any of your FlashSystem A9000R.
- **Site security considerations**
  When installing an IBM FlashSystem A9000R storage system, you need to apply the same security practices that you apply to any other business critical IT system.

Floor and space requirements

Ensure that the location of the FlashSystem A9000R system meets floor and space requirements.

Procedure

Complete the following steps to ensure that the planned installation location meets space and floor load requirements:

1. Decide whether the FlashSystem A9000R system is to be installed on a raised floor.
   See Raised or non-raised floor considerations.
2. Determine whether the floor meets the floor-load requirements for the FlashSystem A9000R system.
   See Floor-load requirements.
3. Calculate the amount of space needed for the rack footprint and service clearance requirements.
   See Rack dimensions and service clearance requirements.
4. Determine where to place the rack in the installation site, based on the floor-load and space requirements.
5. If the location has a raised floor, prepare the raised floor with cable cutouts and required ventilation.
   See Preparing for raised-floor installation and cabling.
6. If the location is not a raised floor, resolve any safety concerns that are caused by the location of overhead-cable exits and cable routing.
   See Preparing for non-raised-floor installation and cabling.
7. Provide your IBM® service representative with the following information before the installation:
   a. Whether under-floor or over-head power-cabling scheme is to be used.
   b. The distance of the rack from the power receptacles.
8. If a rear-door heat exchanger is being ordered, be sure to follow instructions in Preparing for the rear-door heat exchanger.
9. If absorbent padding is used where the rack casters (wheels) are located, be sure to follow instructions in Bottom rack dimensions.

- **Raised or non-raised floor considerations**
  The IBM FlashSystem A9000R storage system can be installed on a raised or a non-raised floor.
- **Floor-load requirements**
  You must ensure that the floor load rating can support the weight of the FlashSystem A9000R system.
- **Rack dimensions and service clearance requirements**
  The installation site must accommodate the rack dimensions and minimum service clearance for the FlashSystem A9000R system.
- **Preparing for raised-floor installation and cabling**
  Prepare the raised floor with cable cutouts, required ventilation, and additional floor support, if necessary.
- **Preparing for non-raised-floor installation and cabling**
  Prepare the installation site to accommodate overhead cabling for mainline-power cables, customer Fibre Channel (FC) and Ethernet host cables as well as network Ethernet cables.
- **Preparing for the rear-door heat exchanger**
  An optional rear-door heat exchanger (feature code AFR1) may be ordered to help cool your system.

Raised or non-raised floor considerations

The IBM FlashSystem® A9000R storage system can be installed on a raised or a non-raised floor.

Raised floor considerations

Installing the racks on a raised floor provides the following benefits:

- Improves operational efficiency and provides greater flexibility in the arrangement of equipment.
- Increases air circulation for better cooling.
- Protects the interconnecting cables and power receptacles.
- Prevents tripping hazards because cables can be routed underneath the raised floor.

When you install on a raised floor, consider the following factors:

- The raised floor must be constructed of fire-resistant or noncombustible material.
- Avoid the exposure of metal or highly conductive material at ground potential to the walking surface when a metallic raised floor structure is used. Such exposure is considered an electrical safety hazard.
- The raised floor height must be at least 30.5 cm (12 in.). Clearance must be adequate to accommodate interconnecting cables, Fibre Channel (FC) cable raceways, power distribution, and any piping that is present under the floor. Floors with greater raised floor heights allow for better equipment cooling.
- When a raised floor tile is cut for cable entry or air supply, an extra floor tile support (pedestal) might be required to restore the structural integrity of the panel to the previous requirement.
- The use of a protective covering (such as plywood, tempered masonite, or plyron) is required to prevent damage to floor tiles, carpeting, and panels while equipment is being moved into or is relocated within the installation site. When the equipment is moved, the dynamic load on the casters is greater than when the equipment is stationary.
- Concrete subfloors require treatment to prevent the release of dust.
- Use noncombustible protective molding to eliminate sharp edges on all floor cutouts to prevent damage to cables and hoses, and to prevent casters from rolling into the floor cutout.
- Seal raised-floor cable openings to prevent the escape of chilled air.
- Pedestals must be firmly attached to the structural (concrete) floor by using an adhesive.

For more information, see Preparing for raised-floor installation and cabling.

Non-raised floor considerations

Raised floors are preferred because they provide better support for the cabling and ensure efficient cooling for the FlashSystem A9000R system; however, overhead cabling at the rear of the rack is available when the FlashSystem A9000R system is installed on a non-raised floor.

Unlike raised-floor cabling, the installation planning, cable length, and the rack location, in relation to the cable opening at the top of the rack, are critical to the successful installation when using overhead cabling.

For more information, see Preparing for non-raised-floor installation and cabling.

Floor-load requirements

You must ensure that the floor load rating can support the weight of the FlashSystem A9000R system.

Floor reinforcement must support the weight of the FlashSystem A9000R system over a specific area, as shown in Table 1. These measurements are slightly less than the footprint area of the system, due to the overhang of the door.

To support future scale-out capability (MES) upgrades, the installation site must provide floor weight-support requirements to support a full rack configuration.

Table 1. Floor weight-support requirements

Grid element configuration	Floor reinforcement area	Total weight
12 grid controllers and 6 flash enclosures Model 415 only.	66 cm × 118 cm (26 in. × 46.5 in.)	935 kg (2061 lbs)
10 grid controllers and 5 flash enclosures Model 415 only.	66 cm × 118 cm (26 in. × 46.5 in.)	855 kg (1885 lbs)
8 grid controllers and 4 flash enclosures	66 cm × 118 cm (26 in. × 46.5 in.)	774 kg (1706 lbs)
6 grid controllers and 3 flash enclosures	66 cm × 118 cm (26 in. × 46.5 in.)	695 kg (1532 lbs)
4 grid controllers and 2 flash enclosures	66 cm × 118 cm (26 in. × 46.5 in.)	616 kg (1358 lbs)
3 grid controllers and 1 flash enclosure	66 cm × 118 cm (26 in. × 46.5 in.)	550 kg (1212 lbs)

Note: The same floor-loading requirements apply to partially populated and fully populated racks.
The rear-door heat exchanger adds more weight to the rack. For information about the weight of an empty and filled door, see Rear-door heat exchanger.

To ensure that all requirements are met, obtain the service of a qualified structural engineer to prepare the floor.
Attention: If you do not know or are not certain about the floor-load rating of the installation site, you must check with the building engineer or another appropriate person.

Rack dimensions and service clearance requirements

The installation site must accommodate the rack dimensions and minimum service clearance for the FlashSystem A9000R system.

The IBM® service representative must have enough space to open the front and rear covers to service the FlashSystem A9000R system, including removing components and other assemblies from the FlashSystem A9000R system.
Note:

- You can position racks no closer than 45 cm (17.7 in.) to a wall.
- You can position racks alongside (next to) other racks.
- Because several rack designs are available from IBM and other vendors, space between adjacent racks might be required to open the door for service. You must determine the space requirement at the time of installation.

Table 1 and Figure 1 describe dimensions and minimum service clearance for the FlashSystem A9000R system.

Table 1. Rack dimensions and clearance requirements

Dimension	Clearance
Height	201.5 cm (79.3 in.)
Depth	129.7 cm (51.1 in.)
Width	64.4 cm (25.4 in.)
Front clearance	120 cm (47.2 in.)
Rear clearance	100 cm (39.4 in.)
Side (door) clearance	10 cm (3.9 in.)

Figure 1. Clearance requirements for servicing the FlashSystem A9000R rack

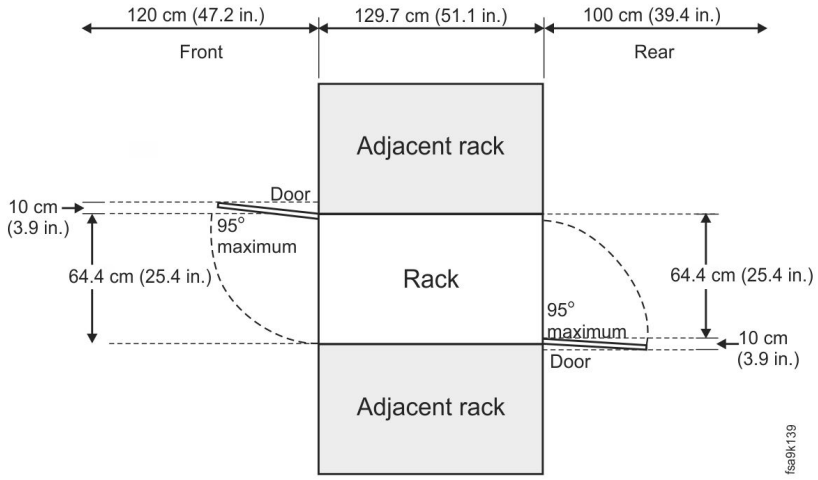

- **Bottom rack dimensions**
  When using absorbent padding where the rack casters (wheels) are located, use this information for proper pad placement.

Bottom rack dimensions

When using absorbent padding where the rack casters (wheels) are located, use this information for proper pad placement.

Figure 1 shows the bottom rack dimensions and castor placements.

Figure 1. Bottom rack dimensions and castor placements

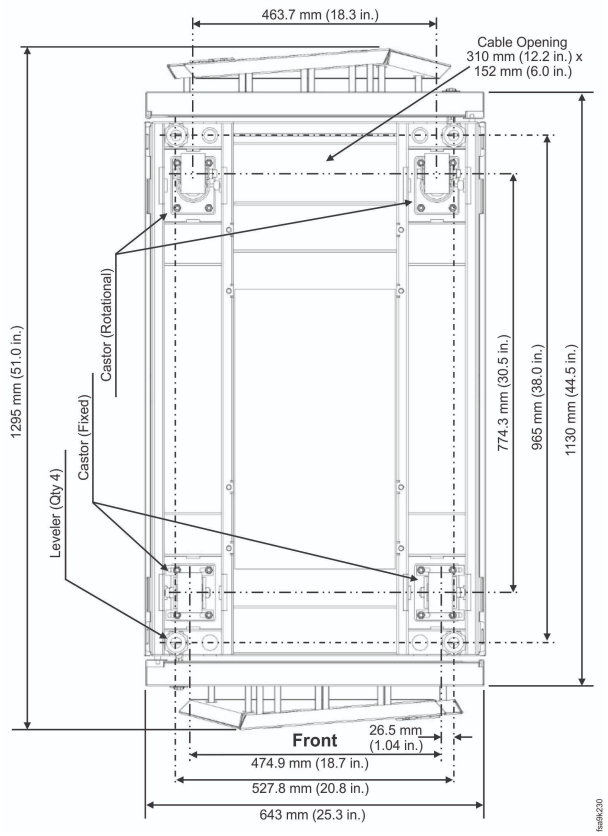

Preparing for raised-floor installation and cabling

Prepare the raised floor with cable cutouts, required ventilation, and additional floor support, if necessary.

Procedure

Complete the following steps to prepare for cabling each FlashSystem A9000R based on raised floor with 60 × 60 cm (24 × 24 in.) tiles:

1. Based on your planned layout, ensure that the installation site can accommodate the locations of the cables exiting each FlashSystem A9000R rack.
2. Plan for the FlashSystem A9000R system to be positioned on two tiles, with the rear of the rack aligned on a floor-tile seam and with two full rows of perforated tiles immediately in front of the rack (see Figure 1).

Figure 1. Raised floor requirements

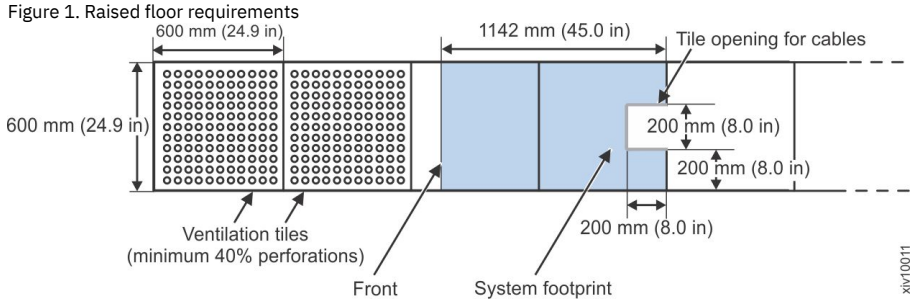

3. Cut a 200 × 200 mm (8 × 8 in.) opening in the rear floor tile for under-floor cabling and electricity (see Figure 1).
   Place the opening under the FlashSystem A9000R system, centered on the back edge of the tile along the rear of the rack.
   Important: Be sure to size the cutout correctly. An oversized cutout permits excessive cooling loss and weakens the floor tile. An undersized cutout must be enlarged, which causes an installation delay while the tile is replaced or the cutout is enlarged.
4. To allow for ventilation for airflow and support system cooling requirements, have at least two tiles (and preferably more) in front of the FlashSystem A9000R system (see Figure 1). These tiles must have a minimum of 40% perforation.
5. If the rear-door heat exchanger is ordered, see Raised floor hose requirements and management for additional floor preparation steps.

Preparing for non-raised-floor installation and cabling

Prepare the installation site to accommodate overhead cabling for mainline-power cables, customer Fibre Channel (FC) and Ethernet host cables as well as network Ethernet cables.

Using overhead cabling provides many of the cooling and safety benefits that are provided by raised flooring in a non-raised floor environment.

Unlike raised-floor cabling, the installation planning, cable length, and the system location in relation to the cable entry point are critical to the successful installation of a top cable exit.

- Main power cables are routed to the rack by the customer, and internally routed and connected by an IBM® service representative.
- Host-attachment cables are internally routed and connected by either the customer or by an IBM service representative.
- All remaining cables are internally routed and connected by an IBM service representative.

If the rear-door heat exchanger is ordered, see Non-raised floor hose requirements and management for more floor preparation steps.

Installation and safety requirements

If the cables are too long, there might not be enough room inside of the rack to handle the extra length and the extra cable might interfere with the servicing tasks, preventing concurrent repair.

IBM Corporate Safety restricts the servicing of your overhead equipment to a maximum of 10 feet from the floor. Therefore, your power source must not exceed 10 feet from the floor and must be within 5 feet of the top of the rack.

Servicing any overhead equipment higher than 10 feet requires a special bid contract. Contact your IBM Representative for more information about special bids.

Preparing for the rear-door heat exchanger

An optional rear-door heat exchanger (feature code AFR1) may be ordered to help cool your system.

To complete the rear-door heat exchanger site preparation, follow the instructions in Planning for the rear-door heat exchanger.

Power requirements

Ensure that your operating environment meets the AC-power and voltage requirements for IBM FlashSystem® A9000R systems.

The FlashSystem A9000R system is designed with backup battery modules in order to maintain power to the storage system in the event of an AC-power loss.

The FlashSystem A9000R system has redundant main power cables. For two-main-power-cable configuration, you must supply power from two independent sources of electricity.

Consult with an IBM® service representative to discuss power source options for the four-main-power-cable configuration.
Note: Removing all AC power from the FlashSystem A9000R system causes an emergency shutdown. All modified data is then saved to drives, and the system turns off within 5 minutes.

Customer responsibilities

The following are the customer responsibilities for ensuring that your operating environment meets all power requirements.

- You must supply enough branch circuits to prevent overloading from the equipment that you install.
  At least two separate power grids are necessary for each system.
- You must ensure that each electrical outlet is correctly wired and grounded to prevent an electrical shock.

IBM responsibilities

The following are the IBM responsibilities for ensuring that your operating environment meets all power requirements.

- The IBM service representative completes several checks, including voltage and grounding checks before the power to the FlashSystem A9000R system is connected.
- The IBM service representative connects power to the racks and initially powers on the equipment.

Power outlet requirements

Ensure that the installation site has the required power outlets.

Two independent power grids are required for the main power cords for each FlashSystem A9000R system.

For systems with single-phase PDUs, two separate outlets on each of the power grids are required for the main power cords.

In order to eliminate a single point of failure, the main power cords from each PDU must connect to separate power grids (sources) and each power grid must have its own wall circuit breaker, as shown in Figure 1.

Figure 1. FlashSystem A9000R power sources

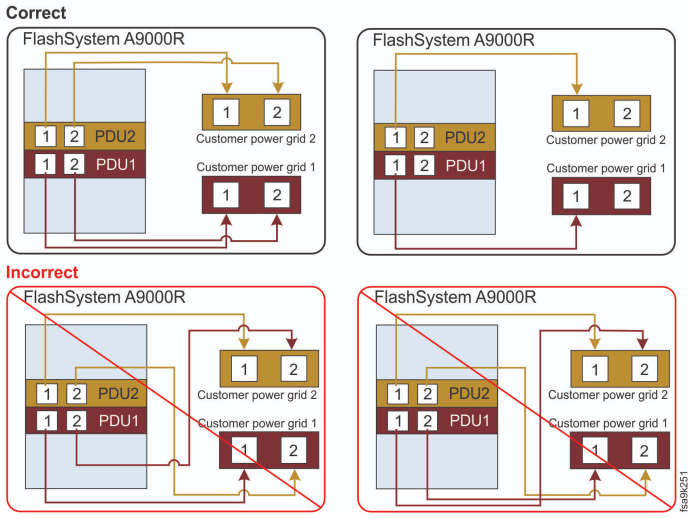

Note: When systems only require one main-power cable per power distribution unit (PDU), only Input 1 from each PDU is used. Two independent power grids are still required.
For the most reliable operation, do not use Ground Fault Circuit Interrupter (GFCI), Earth Leakage Circuit Breaker (ELCB), and Residual Current Circuit Breaker (RCCB) type circuit breakers with the FlashSystem A9000R system.

The FlashSystem A9000R system is certified for safe operation and is compliant with IEC, EN, UL, CSA 60950-1 standards. However, if leakage detection circuit breakers are required by local electrical practice, the breakers must be sized for a leakage-current rating of 100 mA or greater to reduce the risk of server outage caused by erroneous and spurious tripping.

- **Power sources**
  Several AC power source configurations are available.
- **Power consumption**
  This information describes the power consumption for partial and full rack configurations.
- **Input voltages and frequencies**
  This information lists the voltages or frequencies that are accepted by the FlashSystem A9000R system.
- **Main power cables specifications**
  Plan for power cable, connector, and receptacle requirements.
- **Main power cables considerations**
  Use this information when considering the different main power cabling options.

Power sources

Several AC power source configurations are available.

- (Model 415 only) Four 60/63 A, 200-240 V AC, North American, EMEA, and Japan single-phase receptacles, each connected to a different power source.
- (Models 425 and U25 only) Two or four (dependent on configuration's scaled-out power load) 60/63 A, 200-240 V AC, North American, EMEA, and Japan single-phase receptacles, each connected to a different power source.
- Two 60 A, 200-240 V AC, US and Japan delta three-phase receptacles, each connected to a different power source.
- Two 30/32 A, 200-240 V AC (Line-to-Neutral [LN]), EMEA WYE three-phase receptacles, each connected to a different power source

The storage system is protected from a power outage by internal backup battery modules. However, you can reduce the risk of a power outage by connecting the system to an external uninterruptible power supply, a backup generator, or both.

Power consumption

This information describes the power consumption for partial and full rack configurations.

Table 1 and Table 2 list the power consumption for each rack configuration.
The power consumption for partial-rack configurations is the typical value and is provided only for reference purposes.

To support concurrent capacity upgrades (scale-out procedures), the installation site must provide sufficient power capacity to support performance on the new rack configuration.
Note: In order for the installation site to support any future capacity upgrade procedures, the site should adhere to all requirements for a full rack system.
The measurements in Table 1 and Table 2 were taken in an environment with a room temperature of 18°C (64.4°F), all fans at nominal/idle speed, and battery modules were not charging.

Table 1. Power consumption – model 415

Configuration	2.9 TB MicroLatency® modules[1] Idle / light load power consumption		5.7 TB MicroLatency modules Idle / light load power consumption	
	kVA	kW	kVA	kW
4 grid controllers and 2 flash enclosures	3.420	3.089	3.820	3.450
6 grid controllers and 3 flash enclosures	5.034	4.547	5.634	5.089
8 grid controllers and 4 flash enclosures	6.448	5.824	7.448	6.728
10 grid controllers and 5 flash enclosures	8.262	7.363	9.262	8.366
12 grid controllers and 6 flash enclosures	9.876	8.921	11.076	10.005
[1]The numbers for power consumption of 2.9 TB MicroLatency modules are calculated estimates.				

Table 2. Power consumption – models 425 and U25

Configuration	3.6 TB MicroLatency modules[1] Idle / light load power consumption		8.5 TB MicroLatency modules[1] Idle / light load power consumption		18 TB MicroLatency modules Idle / light load power consumption	
	kVA	kW	kVA	kW	kVA	kW
3 grid controllers and 1 flash enclosure	2.027	1.834	2.071	1.874	2.115	1.914
4 grid controllers and 2 flash enclosures	3.258	2.940	3.338	3.012	3.418	3.085
6 grid controllers and 3 flash enclosures	4.791	4.324	4.911	4.432	5.031	4.540
8 grid controllers and 4 flash enclosures	6.324	5.707	6.484	5.852	6.644	5.996
[1]The numbers for power consumption of 3.6 TB and 8.5 TB MicroLatency modules are calculated estimates.						

Input voltages and frequencies

This information lists the voltages or frequencies that are accepted by the FlashSystem A9000R system.

Table 1 provides the voltages and frequencies that are accepted by the FlashSystem A9000R system. For more information, see Leakage current.

Table 1. Input voltages and frequencies

Characteristics	3ØΔ (3-phase Delta) (3+PE) voltage or frequency	3ØY (3-phase Wye) (3+N+PE) voltage or frequency	1Ø (Single-phase, Line-to-Neutral (LN) or Line-to-Line (LL))
Nominal input operating rated range voltages	200 - 240 V AC	LL: 346 - 415 V AC LN: 200 - 240 V AC	200 - 240 V AC
Maximum branch circuit protection rating	60/63 A	30/32 A	60/63 A
Minimum tolerant input voltage	180 V AC	LL: 312 V AC LN: 180 V AC	180 V AC
Maximum tolerant input voltage	264 V AC	LL: 457 V AC LN: 264 V AC	264 V AC
Maximum ground-leakage current	3.8 mA	5.1 mA	11.6 mA
Steady-state input frequencies	50±3 Hz or 60±3 Hz	50±3 Hz or 60±3 Hz	50±3 Hz or 60±3 Hz
Power line disturbance (PLD) input frequencies	50±3 Hz or 60±3 Hz	50±3 Hz or 60±3 Hz	50±3 Hz or 60±3 Hz

Main power cables specifications

Plan for power cable, connector, and receptacle requirements.

The FlashSystem A9000R system is available in multiple main power cable configurations.

Attention: The regional designations are general. If the input voltage for the country uses a wye circuit, use the appropriate main power cables for EMEA (Europe, Middle East, and Africa) and Asia/Pacific.
If input voltage for the country uses a delta circuit, use the appropriate main power cables for United States, Canada, Latin America.

For more information about electric currents for various countries, see the International Trade Administration website .

Main power cable feature codes

The following table provides the feature codes and description details for each main power cable.

- The main power cables are provided with the connector feature codes listed in the following table. Receptacle feature codes in the following table are recommended.
- Although equivalent receptacles can be used, it is the responsibility of the customer to verify compatibility.
- The IEC 60309 standard commercial/industrial pin and sleeve power connectors are often abbreviated "IEC '309" or simply "309 wall plug".

Table 1. Main power cables

Main-power cable feature code	Wall connector	Description	Line cord connector	Customer facility	Inline receptacle	System (PDU) connector
(Model 415) 1050 (Models 425 and U25) 1054	Plug: 60/63 A 250VAC-IEC309-2P3W-IP67	Main-power cable with IEC'309 service ready plug, for: United States, Canada, Latin America, Japan , and Taiwan (EMEA and other regions can optionally order); single-phase, 60 A plug-class (or 60 A rated class plug), two conductors+ground with plug 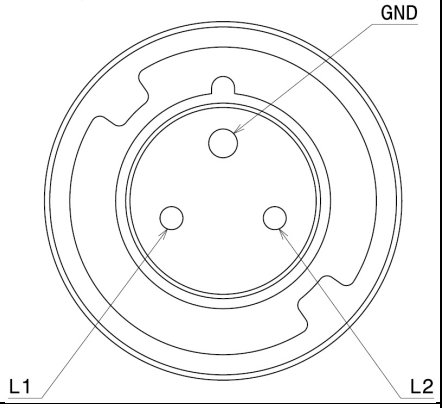	Hubbell HBL360P6V04	Hubbell HBL360R6W	Hubbell HBL360C6W	Single-phase PDU 3- position appliance coupler, application specific
(Model 415) 1051 (Models 425 and U25) 1055	N/A (stripped wires)	Main-power cable with no customer connector provided (pre-stripped wires, termination ready for EMEA (except Japan), hard-wired, or uses client provided customer connector/plug as regionally allowed); single-phase, 60/63 A rated branch-circuit class, two conductors+ground, hardwired 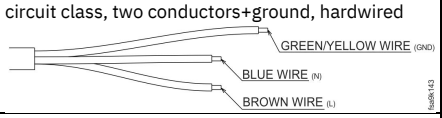	Hardwired or uses client-provided connector	Hardwired AC supply	N/A	Single-phase PDU 3- position appliance coupler, application specific
1052	Plug: 60/63 A 250VAC-IEC309-3P4W-IP67	Main power cable with IEC '309 service ready plug, for: United States, Canada, Latin America, Japan, and Taiwan (EMEA and other regions can optionally order); three-phase delta, 60 A rated plug-class, three wires+ground with plug 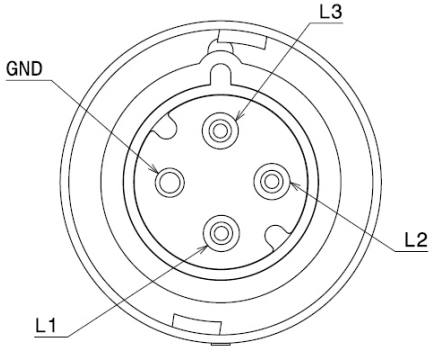	Hubbell HBL460P9V05	Hubbell HBL460R9W	Hubbell HBL460C9W	Delta PDU 4-position appliance coupler, application specific

Main-power cable feature code	Wall connector	Description	Line cord connector	Customer facility	Inline receptacle	System (PDU) connector
1053	N/A (stripped wires)	Main power cable with no customer connector provided (pre-stripped wires, termination ready for EMEA (except Japan), hard-wired, or uses client provided customer connector/plug as regionally allowed) for Europe, Middle East, Asia/Pacific (except Japan), Australia, and New Zealand; three-phase wye, 30/32 A rated branch-circuit class, three wires+neutral + ground, hardwired 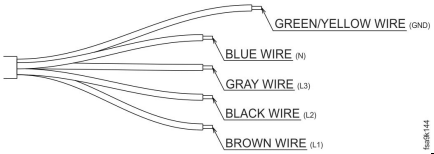	Hardwired or uses client-provided connector	Hardwired AC supply	N/A	Wye PDU 5-position appliance coupler, application specific
1056 Models 425 and U25 only.	Plug: 30/32 A	Main power cable with IEC'309 service ready plug, for EMEA (other regions under approved circumstances, such as United States, Canada, Latin America, Japan, and Taiwan can optionally order); three-phase wye, 30/32 A rated branch-circuit class, three wires+neutral+ground with plug 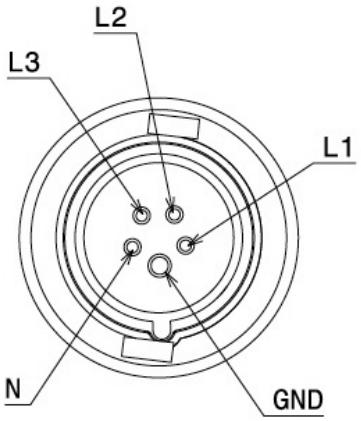	Hubbell HBL530P6-V02	Hubbell HBL530P6-V02	Hubbell HBL530P6-V02	Wye PDU 5-position appliance coupler, application specific

Main power cables considerations

Use this information when considering the different main power cabling options.

The following information provides details to consider regarding main power cables for IBM FlashSystem® A9000R systems:

- Main power cables are up to 250 V 60/63 A or 250 V 30/32 A rated maximum branch circuit classes.
  Single-phase main power cables have two conductors (where the second conductor may be referred to as L2 or N, depending on the installation region mains type implementation) and three wires.

  Three-phase main power cables have three poles and four wires, or four poles and five wires.

  Plugs and receptacles for feature codes 1050, 1052, 1054, 1056 are IEC309-compliant.

- Bulk wire conductor sizes for main power cable feature codes are:
  - 1050, 1051, 1052, 1054, and 1055: 8 AWG
  - 1053, 1056: 6 mm^2
- The main power cables extend 4.1 m (13 ft 4 in.) when they exit from the top of the frame and 4.3 m (14 ft) when they exit from the bottom of the frame.
- Each feature code has specific wall circuit breaker requirements. Choose the best powering options according to your circuit breaker needs:
  - Wall circuit breakers that have a rating of 60 A to 63 A require single-phase installations (200 V - 240 V) with two main power cables.
    For more information regarding single-phase minimum rack installations for models 425 and U25, see Special considerations when powering a rack containing up to four grid controllers with a single-phase power distribution unit (models 425 and U25).

  - Wall circuit breakers that have a rating of 30 A to 60 A require three-phase delta installations (200 V - 240 V).
  - Wall circuit breakers that have a rating of 30 A to 32 A require three-phase wye installations (220 V - 240 V (Line-to-Neutral [LN])).

  Attention: Do not exceed the wire rating of the facility. Wall circuit breakers must be appropriate for power-supply applications to allow internal circuit breakers to function correctly and to avoid unnecessary service.
  For the most reliable operation, do not use Ground Fault Circuit Interrupter (GFCI), Earth Leakage Circuit Breaker (ELCB), and Residual Current Circuit Breaker (RCCB) type circuit breakers with the FlashSystem A9000R system.

  The storage system is certified for safe operation and is compliant with IEC, EN, UL, CSA 60950-1 standards. However, if leakage detection circuit breakers are required by local electrical practice, the breakers must be sized for a leakage-current rating of 100 mA or greater to reduce the risk of server outage caused by erroneous and spurious tripping. For more information, see fs9kx_sftleakagecurrent.html.

- Main power cable cords (bulk wire) of the above listed feature codes have met universal combination certifications: UL/CSA bi-national certification for North America as well as EU Harmonization (HAR) European approvals.
  Other regional approvals might be added as necessitated for industrial/commercial implementations.

Special considerations when powering a rack containing up to four grid controllers with a single-phase power distribution unit (models 425 and U25)

Note: These considerations apply for all model 425 and U25 racks with single-phase PDUs containing up to four grid controllers and two flash enclosures.
When powering a FlashSystem A9000R system with single-phase power distribution units and three or more grid elements, both main input cables of each PDU must always be connected.
When powering a FlashSystem A9000R model 425 or U25 less than three grid elements containing single-phase PDUs, power is taken only from Input 1 from each PDU. Therefore there are different main power cabling options.

For more information regarding the various types of rack configurations, see Rack configurations.
Note: All single-phase PDUs come with two main power cables, no matter the rack configuration.
There are two main power cabling scenarios with single-phase PDUs for racks containing four or less grid controllers. Choose the one that best suits your needs, together with your IBM representative.

1. Both main power cables are connected to the PDU and to the power grid wall socket.
   In this scenario, the system will still only draw power from one of the power cables from each PDU.

2. Only one main power cable is connected to each PDU and to each power grid wall socket.
   In this scenario, the second main power cable of each PDU is given to you for safe keeping. These cables are required to support future scale-out capability (MES) upgrades.

Important: Be sure that both your IBM representative **and** the IBM service provider installing your system are informed as to which option you have chosen.
Note: Scale-out (MES) for model 425 and U25 configurations from minimum rack (three grid controllers and one flash enclosure) is supported from system version 12.3.1. If needed, contact your IBM representative for more information.

Environmental requirements

Ensure that the installation site meets all operating environment requirements for the IBM FlashSystem® A9000R system.

Procedure

To ensure that the installation site meets the requirements, complete the following steps:

1. Use adequate ventilation, especially during the first 120 days of continuous operation. Ensure that there is an average room outdoor intake air rate of 0.4 air change per hour.
2. Keep the front and rear of the rack clear of obstruction.
3. Verify that you can meet the environmental operating requirements at the air intake locations.
4. Consider optimizing the air circulation and cooling for the rack by using a raised floor, adjusting the floor layout, and adding perforated tiles around the air intake areas.

- **Operating and shipping environment requirements**
  You must verify that your operating environment is compatible with the required specifications.
- **Air circulation and cooling**
  You can take steps to optimize the air circulation and cooling for your IBM FlashSystem A9000R system.
- **Contamination information**
  You must consider the air quality and contamination levels at your installation site.
- **Acoustic declaration**
  This information lists the acoustic (sound power) levels for the FlashSystem A9000R system.
- **Operating vibration requirements**
  The vibration levels that are designed for the FlashSystem A9000R system comply with class V1L requirements included in the product classes for vibration.

Operating and shipping environment requirements

You must verify that your operating environment is compatible with the required specifications.

With or without the rear-door heat exchanger option, the air that enters the front door of the FlashSystem A9000R system must meet the following requirements.

Operating (powered on)

- Temperature: 10°C to 35°C (50°F to 95°F)
- Relative humidity: 20% to 80%, non-condensing
- Maximum wet bulb temperature: 25°C (77°F)
- Maximum altitude: 2134 m (7000 ft)

Non-operating (powered off)

- Temperature: 5°C to 45°C (41°F to 113°F)
- Relative humidity: 20% to 80%, non-condensing

- Maximum wet bulb temperature: 27°C (80.6°F)

Shipping

- Temperature: -40°C to 60°C (-40°F to 140°F)
- Relative humidity: 5% to 95%, non-condensing
- Maximum wet bulb temperature: 29°C (84.2°F)

Important: A FlashSystem A9000R system that runs continuously must be within the specified operating environment.
A storage system can operate at the maximum allowable temperature for only short durations, such as might occur during a disk drive module or power-supply unit replacement.

Continuous operation above the maximum temperature increases the probability of component failure.

Air enters at the front of the rack and leaves at the back. To prevent the air that is leaving the rack from entering the intake of another piece of equipment, place racks in alternate rows, in a back-to-back, and front-to-front arrangement.
Align the front of racks on a floor-tile seam, with a full line of perforated tiles immediately in front of the rack.

- **System environmental acclimation**
  Condensation is a normal and natural occurrence if packaging is removed at time of delivery, during extreme cold weather delivery times.

System environmental acclimation

Condensation is a normal and natural occurrence if packaging is removed at time of delivery, during extreme cold weather delivery times.

All IBM® equipment has been tested in climate chambers duplicating shipping temperatures of -40°C (-40°F) to 60°C (140°F) and condensation and frost do not impact reliability of the product.

It is highly recommended that the device not be removed from the shipping package for at least 24 hours, in order to acclimate to the new temperature conditions. If there are still visible signs of condensation after the initial 24 hour wait period, acclimate the system without the shipping bag for an additional 12 - 24 hours, or until no further visible condensation remains. If package material removal is necessary, equipment should be placed in a location away from any air vents.

Air circulation and cooling

You can take steps to optimize the air circulation and cooling for your IBM FlashSystem® A9000R system.

Procedure

To optimize the cooling around your FlashSystem A9000R system, complete the following steps:

1. Install the FlashSystem A9000R system on a raised floor, which provides increased air circulation for better cooling.
   For more information, see Raised or non-raised floor considerations.

2. Install perforated tiles in the front and back of each base rack and expansion rack as follows:
   a. For a stand-alone base rack, install two fully perforated tiles in front of the base rack and one partially perforated tile at the back of the base rack.
   b. For a row of racks, install a row of perforated tiles in front of the racks and one or two fully perforated tiles at the back of each two racks.
   c. For groupings of racks, where a hot aisle and cold aisle layout is used, use a cold aisle row of perforated tiles in front of all racks. For hot aisles, install a perforated tile per pair of racks.
   For more information, see Preparing for non-raised-floor installation and cabling.

3. Consider using the rear-door heat exchanger (feature code AFR1) to reduce the temperature of the air that leaves the rack.
   The rear-door heat exchanger is a water-cooled door that is designed to remove heat that is generated from the rack before it enters the computer room. The door uses standard fittings and couplings. The door removes up to 50,000 BTUs of heat, which is approximately 15 kW, from the air that exits a rack that is full of servers.

   The rear-door heat exchanger can increase server density without increasing cooling requirements, making the option a more cost-effective solution than adding an air-conditioning unit. This door is an effective solution for a data center at the limit of its cooling capacity, but that still has usable floor space to add racks. The rear-door heat exchanger is also an efficient way to deal with computer room hotspots.

   For more information, see Rear-door heat exchanger.

4. Ensure that the installation site meets the cooling (thermal dissipation) requirements that are listed in one of the following:
   Note: To support future scale-out capability (MES) upgrades, the installation site must provide cooling arrangements to support a full rack configuration.
   - **For model 415:** Table 1
   - **For models 425 and U25:** Table 2

Table 1. Thermal dissipation for FlashSystem A9000R system – model 415

Storage configuration	Thermal Dissipation kBTU/hour
4 grid controllers and 2 flash enclosures	16.4
6 grid controllers and 3 flash enclosures	24.14
8 grid controllers and 4 flash enclosures	31.96
10 grid controllers and 5 flash enclosures	39.6

Storage configuration	Thermal Dissipation kBTU/hour
12 grid controllers and 6 flash enclosures	47.3

Table 2. Thermal dissipation for FlashSystem A9000R system – models 425 and U25

Storage configuration	Thermal Dissipation kBTU/hour
3 grid controllers and 1 flash enclosure	7.5
4 grid controllers and 2 flash enclosures	10.5
6 grid controllers and 3 flash enclosures	15.8
8 grid controllers and 4 flash enclosures	21.0

5. Ensure that the installation site meets the airflow requirements that are listed in one of the following:
   Note: To support future scale-out capability (MES) upgrades, the installation site must provide cooling arrangements to support a full rack configuration.
   - **For model 415:** Table 3
   - **For models 425 and U25:** Table 4

Table 3. Airflow requirements for FlashSystem A9000R system – model 415

Storage configuration	Cubic feet per minute (CFM)	
	Nominal temperature (23°C) (73.4°F)	Maximum temperature (35°C) (95°F)
4 grid controllers and 2 flash enclosures	530	970
6 grid controllers and 3 flash enclosures	730	1330
8 grid controllers and 4 flash enclosures	940	1680
10 grid controllers and 5 flash enclosures	1140	2030
12 grid controllers and 6 flash enclosures	1340	2390

Table 4. Airflow requirements for FlashSystem A9000R system – models 425 and U25

Storage configuration	Cubic feet per minute (CFM)	
	Nominal temperature (23°C) (73.4°F)	Maximum temperature (35°C) (95°F)
3 grid controllers and 1 flash enclosure	300	550
4 grid controllers and 2 flash enclosures	410	750
6 grid controllers and 3 flash enclosures	590	1100
8 grid controllers and 4 flash enclosures	770	1450

- **Temperature threshold and events**
  The storage system handles overheating by informing the administrator through warning events and initiating an automatic thermal shutdown as a last resort.

Temperature threshold and events

The storage system handles overheating by informing the administrator through warning events and initiating an automatic thermal shutdown as a last resort.

These event notifications indicate to the administrators if the system temperature is: normal, high, too high, and critically high; above which shutdown will be applied immediately. Table 1 specifies the temperature thresholds and events received.
After a manual or automatic thermal shutdown due to thermal conditions, IBM® support must be immediately contacted. Do not attempt to power up the system before contacting IBM support.

Table 1. FlashSystem A9000R system temperature thresholds and events

Temperature threshold	Event	Description
27°C (80°F)	SYSTEM_TEMPERATURE_IS_OK_NOW	No action required. This event is generated only when the normal system temperature is recovered from a higher temperature.
28°C (82.4°F)	SYSTEM_TEMPERATURE_IS_ABOVE_NORMAL	System temperature is above normal temperature range.
30°C (86°F)	SYSTEM_TEMPERATURE_IS_HIGH	System should be closely monitored, and action to cool down the system is recommended.
32°C (89.6°F)	SYSTEM_TEMPERATURE_IS_TOO_HIGH	Actions to cool down the system must be taken immediately.
35°C (95°F)	SYSTEM_TEMPERATURE_IS_CRITICALLY_HIGH	Manually shutdown the system by using the **shutdown** command.
38°C (100.4°F)	SYSTEM_TEMPERATURE_IS_CRITICALLY_HIGH_SHUTTING_DOWN	Automatic system shutdown is in progress.

Contamination information

You must consider the air quality and contamination levels at your installation site.

Airborne particulates (including metal flakes or particles) and reactive gases acting alone or in combination with other environmental factors, such as humidity or temperature, might pose a risk to the FlashSystem A9000R system hardware. Risks that are posed by the presence of excessive particulate levels or

concentrations of harmful gases include damage that might cause the FlashSystem A9000R system to malfunction or cease functioning altogether. This specification describes limits for particulates and gases that are intended to avoid such damage. The limits must not be viewed or used as definitive limits because numerous other factors, such as temperature or moisture content of the air, can influence the impact of particulates or environmental corrosives and gaseous contaminant transfer.

In the absence of specific limits, implement practices that maintain particulate or gas levels that are consistent with the protection of human health and safety. If IBM® determines that the levels of particulates or gases in your environment damaged the FlashSystem A9000R system, IBM might require implementation of appropriate remedial measures to mitigate such environmental contamination before providing repair or replacement of the FlashSystem A9000R system. Implementation of such remedial measures is a customer responsibility.

The following criteria must be met:

Gaseous contamination

Severity level G1 as per ANSI/ISA 71.04-1985[1], which states that the reactivity rate of copper coupons must be less than 300 Angstroms per month (Å/month, ≈ 0.0039 µg/cm^2-hour weight gain) [2]. In addition, the reactivity rate of silver coupons must be less than 300 Å/month (≈ 0.0035 µg/cm^2-hour weight gain) [3]. The reactive monitoring of gaseous corrosibleness must be conducted approximately 5 cm (2 in.) in front of the rack on the air inlet side at one-quarter and three-quarter frame height off the floor. For environments with special air handling equipment that alters the normal distribution of airflow into the rack, reactive monitoring of gaseous corrosibleness must be conducted at a location on the rack where air is entering at the highest rate.

Particulate contamination

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

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

For data centers with air-side economizers, the choice of filters to achieve ISO class 8 cleanliness depends on the specific conditions present at that data center. The deliquescence relative humidity of the particulate contamination must be more than 60% RH [4]. Data centers must be free of zinc whiskers [5].

[1] ANSI/ISA-71.04.1985. *Environmental conditions for process measurement and control systems: Airborne contaminants.* Instrument Society of America, Research Triangle Park, NC, 1985.

[2] The derivation of the equivalence between the rate of copper corrosion product thickness growth in Å/month and the rate of weight gain assumes that Cu_2S and Cu_2O grow in equal proportions.

[3] The derivation of the equivalence between the rate of silver corrosion product thickness growth in Å/month and the rate of weight gain assumes that Ag_2S is the only corrosion product.

[4] The deliquescence relative humidity of particulate contamination is the relative humidity at which the dust absorbs enough water to become wet and promote corrosion, ion migration, or both.

[5] Surface debris is randomly collected from 10 areas of the data center on a 1.5 cm (0.6 in.) diameter disk of sticky, electrically conductive tape on a metal stub. If examination of the sticky tape in a scanning electron microscope reveals no zinc whiskers, the data center is considered free of zinc whiskers.

Acoustic declaration

This information lists the acoustic (sound power) levels for the FlashSystem A9000R system.

CAUTION:

Depending upon local conditions, the sound pressure might exceed 85 dB(A) during service operations. Hearing protection must be worn when you are in a room that has a FlashSystem A9000R system while either the front or rear door is open or when the front and rear doors are not installed.
Note: Government regulations (such as those prescribed by OSHA or European Community Directives) may govern noise level exposure in the workplace and may apply to you and your server installation. The actual sound pressure levels in your installation depend upon a variety of factors, including the number of racks in the installation; the size, materials, and configuration of the room; the noise levels from other equipment; the room ambient temperature; and employees' location in relation to the equipment. Further, additional factors may be considered when reviewing noise exposures, including, but not limited to, exposure duration and the use of hearing protection. IBM recommends that you consult with a qualified expert in the acoustical noise field prior to operating this system.

Model 415

The acoustic levels are shown in Table 1. These measurements are based on a typical FlashSystem A9000R configuration with 8 grid controllers and 4 flash enclosures. All measurements are in conformance with ISO 7779 and declared in conformance with ISO 9296.
Note:

1. LWAd is the statistical upper-limit A-weighted sound-power level (rounded to the nearest 0.1 B).
2. LpAm is the mean A-weighted emission sound-pressure level that is measured at the 1-meter bystander positions (rounded to the nearest dB).
3. 10 dB (decibel) = 1 B (bel)

Table 1. Acoustic declaration – model 415

	Declared A-weighted sound power level, LWAd (B)	Declared A-weighted sound pressure level, LpAm (dB)
Operating	8.0	62
Idle	8.0	62

Models 425 and U25

The acoustic levels are shown in Table 2. These measurements are based on a typical FlashSystem A9000R configuration with 8 grid controllers and 4 flash enclosures. All measurements are in conformance with ISO 7779 and declared in conformance with ISO 9296.
Note:

1. LWAd is the statistical upper-limit A-weighted sound-power level (rounded to the nearest 0.1 B).
2. LpAm is the mean A-weighted emission sound-pressure level that is measured at the 1-meter bystander positions (rounded to the nearest dB).
3. 10 dB (decibel) = 1 B (bel)

Table 2. Acoustic declaration – models 425 and U25

	Declared A-weighted sound power level, LWAd (B)	**Declared A-weighted sound pressure level, LpAm (dB)**
Operating	7.7	58
Idle	7.7	58

Operating vibration requirements

The vibration levels that are designed for the FlashSystem A9000R system comply with class V1L requirements included in the product classes for vibration.

The FlashSystem A9000R system is designed to operate under the vibration V1L levels that are described in the following tables. More information includes random vibration PSD profile breakpoints and operational shock levels.

Model 415

Table 1. Vibration levels – model 415

Class	***g*rms**	***g* Peak Sine**
V1L	0.10	0.06 @ 50 & 60 Hz
Note: *g* is the peak *g* level of an approximate half-sine pulse.		

Table 2. Random vibration PSD profile breakpoints – model 415

Class	**5 Hz**	**17 Hz**	**500 Hz**
V1L	2.0×10^{-7}	2.2×10^{-5}	2.2×10^{-5}
Note: All values in this table are in g^2/Hz.			

Table 3. Operational shock levels – model 415

Class	**Axis**	***g*[1]**	**pw[2]**
S1	Vertical	3.5	3.0
Note: • *g* is the peak *g* level of an approximate half-sine pulse. • *pw* is the pulse width in milliseconds.			

Models 425 and U25

Table 4. Vibration levels – models 425 and U25

Class	***g*rms**	***g* Peak Sine**
V1H	0.05	0.03 @ 50 & 60 Hz
Note: *g* is the peak *g* level of an approximate half-sine pulse.		

Table 5. Random vibration PSD profile breakpoints – models 425 and U25

Class	**5 Hz**	**17 Hz**	**500 Hz**
V1H	1.0×10^{-7}	5.2×10^{-6}	5.2×10^{-6}
Note: All values in this table are in g^2/Hz.			

Table 6. Operational shock levels – models 425 and U25

Class	**Axis**	***g*[1]**	**pw[2]**
S1	Vertical	3.5	3.0
Note: [1]*g* is the peak *g* level of an approximate half-sine pulse. [2]pw is the pulse width in milliseconds.			

Planning for the rear-door heat exchanger

If the optional rear-door heat exchanger (feature code AFR1) was ordered, you must prepare the installation site before an IBM® service representative can install the rear-door heat exchanger on any of your FlashSystem A9000R.

- **Rear-door heat exchanger operating specifications**
  Rear-door heat exchanger operating specifications provide detailed information for your heat exchanger, including dimensions, weight, air source, water source, water pressure, and water volume.
- **Rear-door heat exchanger performance**
  This information describes the performance of the rear-door heat exchanger.
- **Preparing your site to provide water to the rear-door heat exchanger**
  You must prepare your site to provide water to the rear-door heat exchanger before the rack can be installed.
- **Secondary cooling loop parts and services information**
  IBM supplies a heat exchanger that is designed for IBM enterprise-server racks, a hinge kit (for those racks), and an air purge tool. This information provides sources and information for other parts and services that are needed for correct function and reliability of the secondary cooling loop.
- **Maintenance schedule**
  Maintenance tasks are required at certain time intervals.

Rear-door heat exchanger operating specifications

Rear-door heat exchanger operating specifications provide detailed information for your heat exchanger, including dimensions, weight, air source, water source, water pressure, and water volume.

The following table shows the specifications for the rear-door heat exchanger.

Table 1. Rear-door heat exchanger specifications

Door specifications	Air specifications	Water specifications

Door size
• Depth: 142.6 mm (5.6 in.)
• Height: 1945.4 mm (76.6 in.)
• Width: 639 mm (25.2 in.)
Heat exchanger size
• Depth: 67 mm (2.6 in.)
• Height: 1791.3 mm (70.5 in.)
• Width: 438.6 mm (17.3 in.)
Door assembly weight
• Empty: 29.9 kg (66 lb)
• Filled: 35.6 kg (78.5 lb)
Door heat removal capacity
• For examples of door heat removal capacity, see the illustrations in Rear-door heat exchanger performance.
• In general, the door heat removal capacity percentage increases if one or more of the following events occur:
  ○ The water temperature decreases.
  ○ The water flow increases.
  ○ The server heat loads decrease.
• The door heat removal capacity varies with water temperature, water flow rate, air temperature and flow, and total heat load of the servers; however, a typical high-load cabinet (20 - 32 kW or approximately 70 000 - 105 000 Btu per hour) can achieve 55% - 85% heat removal.
	Air movement
	• Provided by servers and other devices in the rack. No additional air moving devices are required.
	Air source for servers
	• Room air for front of the rack. Air is exhausted from the servers, moves through the heat exchanger, and exits into the room (open loop).
	Air temperature drop
	• With high-heat-load devices, up to 25°C (45°F) between the air that exits the rack devices and the air that exits the heat exchanger.
	Air impedance
	• Air pressure drop across the heat exchanger is equivalent to that of the IBM® acoustic rear door
		Water source
		• User-supplied, compliant with specifications in this document.
		• Couplings on door: 19 mm (0.75 in.)
		• ID hose required: 19 mm (0.75 in.) minimum
		Water pressure
		• Normal operation: <137.93 kPa (20 psi)
		• Maximum: 689.66 kPa (100 psi)
		• Pressure drop across heat exchanger: approximately 48 kPa (7 psi)
		Water volume
		• Exchanger: approximately 2.8 liters (0.75 gallons)
		• Exchanger plus supply and return hoses to the pump unit: Maximum of approximately 15.1 liters (4.0 gallons) excluding pump unit piping and reservoir
		Water temperature
		• If no dew point control:18°C ±1°C (64.4°F ±1.8°F)
		• If the water supply can monitor and adjust the relative-to-room dew point, lower temperature water is allowed.
		Required water flow rate (as measured at the supply entrance to the heat exchanger)
		• Minimum: 22.7 liters (6 gallons) per minute
		• Maximum: 37.9 liters (10 gallons) per minute

Rear-door heat exchanger performance

This information describes the performance of the rear-door heat exchanger.

An example of expected performance of the rear-door heat exchanger is illustrated in Figure 1 for a typical inlet air temperature of 24 °C (75.2 °F), with a fully populated rack near uniform power dissipation, 32 kW heat load, and the node fans running near nominal fan speed (1530 cfm). By selecting the water inlet

temperature and water flow rate, you can estimate the indicated heat removal. These levels can be achieved with normal cable exits from the rack and with a small amount of hot air bypass at the base of the door (small amounts of hot air might escape from the rack without being cooled by the door).

Figure 1. Typical performance of a rear-door heat exchanger, 32 kW heat load . Percentage heat removal as function of water temperature and flow rate. (24°C (75.2 °F) rack inlet air, 32 kW rack load, 1530 cfm air through the rear-door heat exchanger)

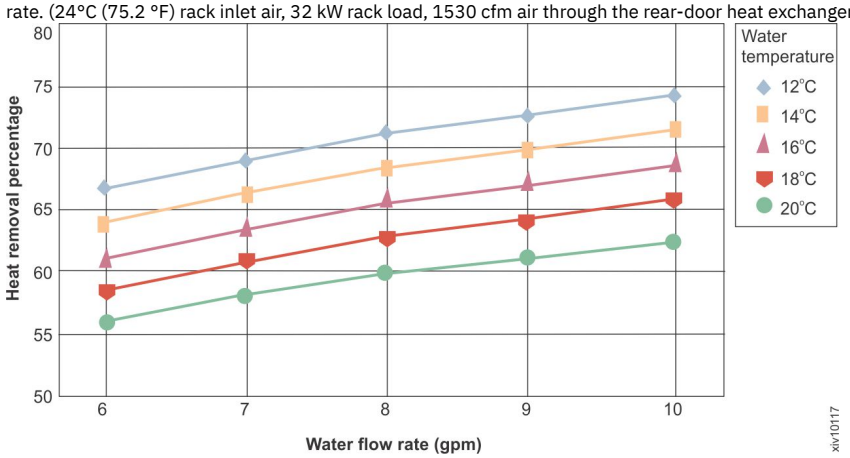

Water temperatures below 18°C (64.4°F) can be used only if the system that is supplying the water is able to measure the room dew point conditions and is able to automatically adjust the water temperature.

Another example of performance data is shown in Figure 2 for identical conditions as in Figure 1, except reflecting a 20 kW heat load. Because of the lower heat load, a specific level of cooling can be achieved with warmer water, a lower flow rate, or both.

Figure 2. Typical performance of a rear-door heat exchanger, 20 kW heat load. Percentage heat removal as function of water temperature and flow rate. (24°C (75.2 °F) rack inlet air, 20 kW rack load, 1530 cfm air through the rear-door heat exchanger)

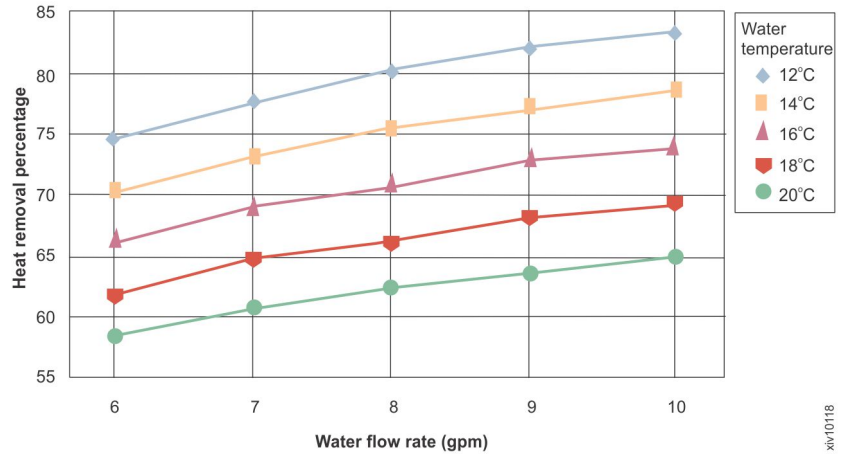

Preparing your site to provide water to the rear-door heat exchanger

You must prepare your site to provide water to the rear-door heat exchanger before the rack can be installed.

The following requirements must be met before the rear-door heat exchanger can be installed:

- Provide chilled, conditioned water that meets the specifications.
- Procure and install the water supply system that is suitable for your data center.
- Provide a redundant secondary-cooling-loop water supply or enough room air conditioning to handle a tolerable heat load if the function of one or more of the heat exchangers is compromised. If the rear door is opened for rack maintenance or if conditioned water supply to the door is stopped, the rack heat load is sent into the room and must be handled by room air conditioning until the conditioned water supply is restored.
- Provide floor or ceiling tile cutouts or protective coverings to avoid tripping hazards on non-raised floors as part of hose management.

- **Water specifications for the secondary cooling loop**
  In order to avoid system failures, it is important to follow the water specification requirements for the secondary cooling loop of your rear-door heat exchanger.
- **Water delivery specifications for secondary loops**
  The delivery system for the secondary cooling loop provides chilled water to the rear-door heat exchanger. The delivery system includes pipes, hoses, and the required connection hardware to connect to the heat exchanger. This information provides examples for setting up the secondary cooling loop and operating characteristics that are needed to provide an adequate, safe supply of water to the heat exchanger.
- **Raised floor hose requirements and management**
  On a raised floor, hoses can be routed under the floor tiles and can be brought up from beneath the rack through special tile cutouts. The hoses are connected to the quick-connect couplings on the bottom of the heat exchanger.
- **Non-raised floor hose requirements and management**
  In data centers without a raised floor, straight hose assemblies cannot make the sharp bend to exit between the floor and the rack door without kinking the

hose.

Water specifications for the secondary cooling loop

In order to avoid system failures, it is important to follow the water specification requirements for the secondary cooling loop of your rear-door heat exchanger.

The water that is being supplied to the heat exchanger must meet the following requirements; otherwise, system failures might occur over time, as a result of:

- Leaks that are caused by corrosion and pitting of the metal components of the heat exchanger or the water supply system.
- Buildup of scale deposits inside the heat exchanger, which can cause the following problems:
  - A reduction of the ability of the heat exchanger to cool the air that is exhausted from the rack.
  - Failure of mechanical hardware, such as a hose quick-connect adapter.
- Organic contamination, such as bacteria, fungi, or algae. This contamination can cause the same problems as described for scale deposits.

Water control and conditioning for the secondary cooling loop

The water that is used to fill, refill, and supply the heat exchanger must be particle-free deionized water or particle-free distilled water with appropriate controls for avoiding the following issues:

- Metal corrosion
- Bacterial fouling
- Scaling

Because of typical water temperatures (described in Preparing your site to provide water to the rear-door heat exchanger), the water might not be able to originate from the primary building chilled-water system. Conditioned water for the heat exchanger must be supplied as part of a secondary, closed-loop system.
Important: Do not use glycol solutions because they can adversely affect the cooling performance of the heat exchanger.

Materials for the secondary cooling loops

You can use any of the following materials in supply lines, connectors, manifolds, pumps, hoses, and any other hardware that makes up the closed-loop water-supply system at your location:

- Copper
- Brass with less than 30% zinc content
- Stainless steel – 303, 304, or 316
- Ethylene Propylene Diene Monomer (EPDM) rubber, peroxide cured, non-metal oxide

Materials to avoid in secondary loops

Do not use any of the following materials in any part of your water supply system.

- Oxidizing biocides (such as, chlorine, bromine, and chlorine dioxide)
- Aluminum
- Brass with greater than 30% zinc
- Irons (non-stainless steel)

Water supply requirements for secondary cooling loops

Ensure that the following requirements are met for the system that supplies the chilled conditioned water to the heat exchanger.

Temperature
: The heat exchanger, supply hose, and return hoses are not insulated and do not have features that are designed to address the creation and collection water from condensate. Avoid any condition that might cause condensation. The temperature of the water inside the supply hose, return hose, and the heat exchanger must be kept above the dew point of the location where the heat exchanger is being used.
Attention: Typical primary chilled water is too cold for use in this application because building chilled water can be as cold as 4°C - 6°C (39°F - 43°F).
Important: If the system that supplies the cooling water is not able to measure the room dew point and automatically adjust the water temperature, the minimum water temperature that must be maintained is 18°C±1°C (64.4°F±1.8°F). The minimum water temperature is consistent with the ASHRAE Class 1 Environmental Specification that requires a maximum dew point of 17°C (62.6°F). See the ASHRAE document Thermal Guidelines for Data Processing Environments. You can find information about obtaining this document on the ASHRAE website .

Pressure
: The water pressure in the secondary loop must be less than 689.66 kPa (100 psi). Normal operating pressure at the heat exchanger must be 137.93 kPa (20 psi) or less.

Flow rate
: The flow rate of the water in the system must be in the range of 23 - 38 liters (6 - 10 gallons) per minute.

  Pressure drop versus flow rate for heat exchangers (including quick-connect couplings) is defined as approximately 48 kPa (7 psi) at 30 liters (8 gallons) per minute.

Water volume limits
: The heat exchangers hold 2.8 - 5.3 liters (0.75 - 1.4 gallons). 15 meters (50 ft) of 19 mm (0.75-in.) supply and return hoses hold approximately 9.4 liters (2.5 gallons). To minimize exposure to flooding in the event of leaks, the entire cooling system (heat exchanger, supply hose, and return hose) excluding any reservoir tank must have a maximum 15.1 liters (4 gallons) of water. This is a cautionary statement not a functional requirement. Also, consider using leak detection methods on the secondary loop that supplies water to the heat exchanger.

Air exposure

The secondary cooling loop is a closed loop, with no continuous exposure to room air. After you fill the loop, remove all air from the loop. Air bleed valves are provided at the top of each heat exchanger manifold for purging all air from the system.

Water delivery specifications for secondary loops

The delivery system for the secondary cooling loop provides chilled water to the rear-door heat exchanger. The delivery system includes pipes, hoses, and the required connection hardware to connect to the heat exchanger. This information provides examples for setting up the secondary cooling loop and operating characteristics that are needed to provide an adequate, safe supply of water to the heat exchanger.

The *primary cooling loop* is considered to be the building chilled-water supply or a modular chiller unit. The primary cooling loop must not be used as a direct source of coolant for the heat exchanger for the following reasons:

- If the supply water temperature is below the room dew point, condensation forms and causes dripping from the door components.
- If a leak develops in the door, supply hose, or return hose, a large amount of water is available.

Procurement and installation of the components that are needed to create the secondary cooling loop system are required for this design and are your responsibility. See Secondary cooling loop parts and services information for information about suppliers of hoses and cooling distribution units.
Attention: The overpressure safety device must meet the following requirements:

- Comply with ISO 4126-1. For more information, go to the ANSI Standards Store website , and search on document number ISO 4126–1.
- Be installed so that it is easily accessed for inspection, maintenance, and repair.
- Be connected as close as possible to the device that it is intended to protect.
- Be adjustable only with the use of a tool.
- Have a discharge opening that is directed so that discharged water or fluid does not create a hazard or directed toward any person.
- Be of adequate discharge capacity to ensure that the maximum working pressure is not exceeded.
- Be installed without a shutoff valve between the overpressure safety device and the protected device.

Figure 4 shows a typical cooling solution and identifies the components of the primary cooling loop and secondary cooling loop.

Figure 3 shows an example of a fabricated facilities solution. The actual number of heat exchangers that are connected to a secondary loop depends on the capacity of the cooling distribution unit that is running the secondary loop.

Figure 1 shows an example of an off-the-shelf modular cooling distribution unit. The actual number of heat exchangers that are connected to a secondary loop depends on the capacity of the cooling distribution unit that is running the secondary loop.

Figure 2 shows an example of a water-chiller unit that supplies conditioned water to one or more heat exchangers. This water-chiller unit must be a closed system (no exposure of the water to air) and meet all materials, water quality, water treatment, and temperature and flow specifications that are defined in this document. A water chiller unit is considered an acceptable alternative to use as a building chilled water source for removing heat from the rear-door heat exchanger.

Figure 1. Cooling distribution unit that uses off-the-shelf supplier solutions

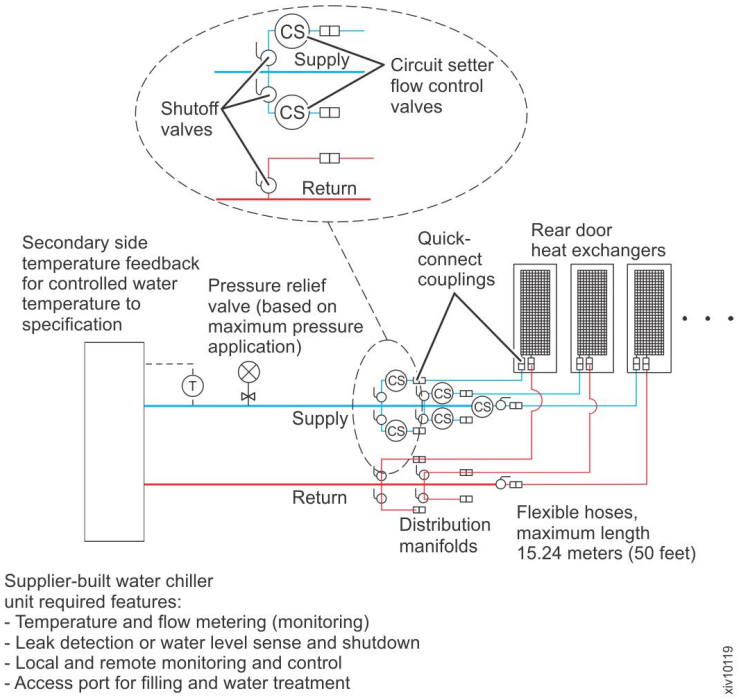

Figure 2. Cooling distribution unit that uses a water chiller unit to provide conditioned water

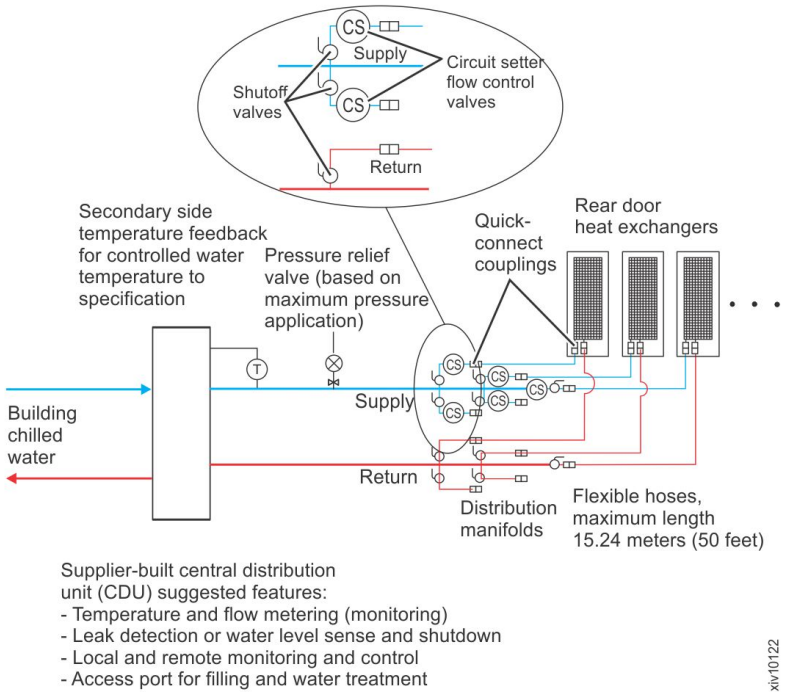

Figure 3. Cooling distribution unit that uses a fabricated facilities solution

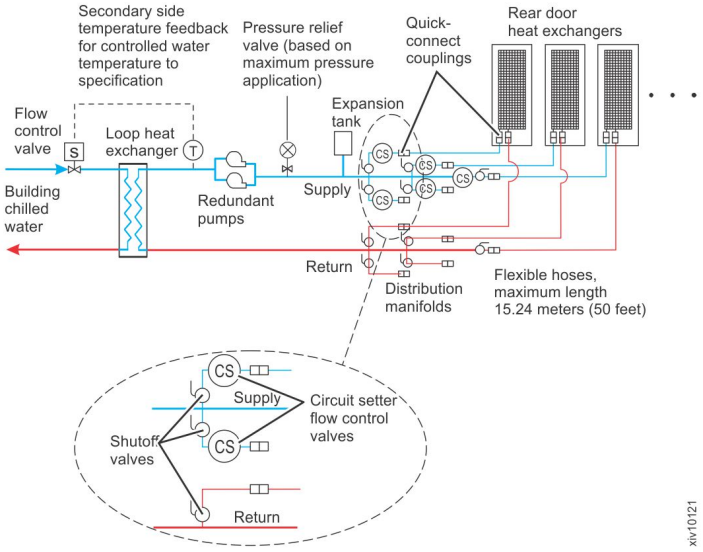

Figure 4. Primary and secondary cooling loops

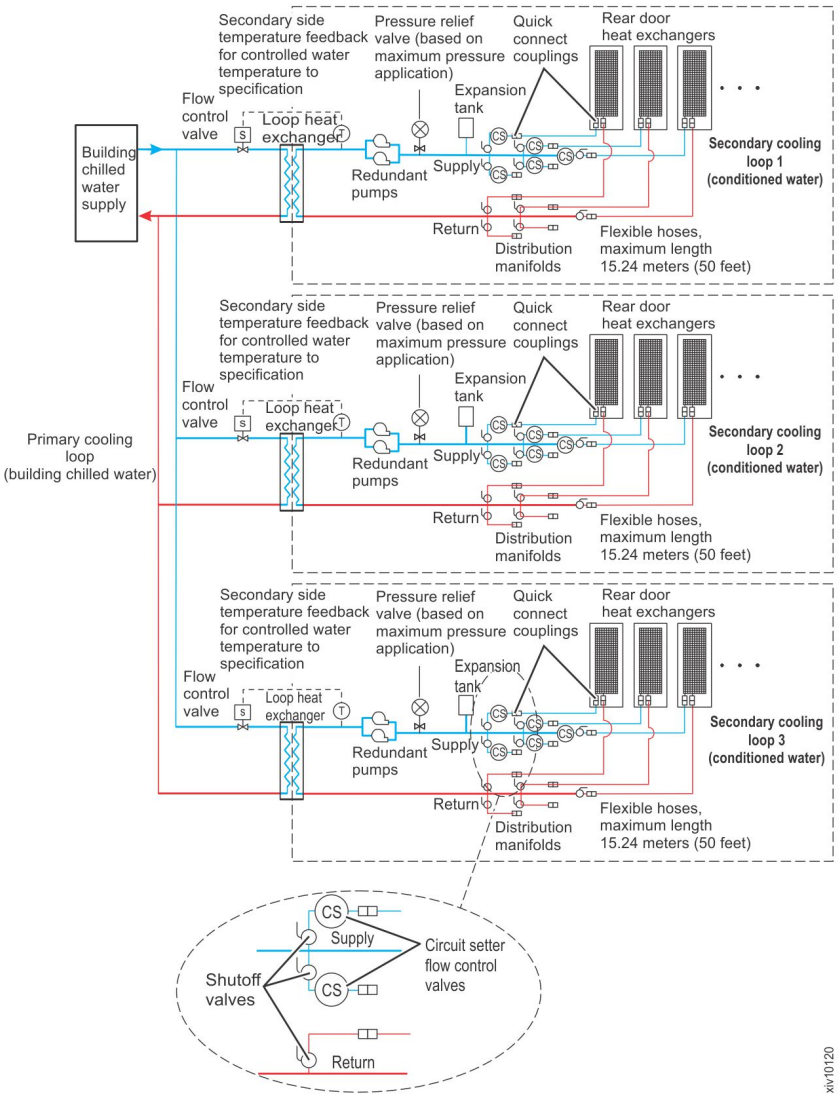

Manifolds and piping

Manifolds that accept large-diameter feed pipes from a pump unit are the preferred method for splitting the flow of water to smaller-diameter pipes or hoses that are routed to individual heat exchangers. Manifolds must be constructed of materials that are compatible with the pump unit and related piping (see Water specifications for the secondary cooling loop). The manifolds must provide enough connection points to allow a matching number of supply and return lines to be attached, and the manifolds must match the capacity rating of the pumps and the loop heat exchanger (between the secondary cooling loop and the building chilled-water source). Anchor or restrain all manifolds to provide the required support to avoid movement when quick-connect couplings are connected to the manifolds.

Manifold supply pipe sizes

- Use a 50.8 mm (2 in.) supply pipe to provide the correct flow to six 19 mm (0.75 in.) supply hoses, with a 100 kW cooling distribution unit.
- Use a 63.5 mm (2.50 in.) supply pipe to provide the correct flow to eight 19 mm (0.75 in.) supply hoses, with a 120 kW CDU.
- Use an 88.9 mm (3.50 in.) supply pipe to provide the correct flow to twenty 19 mm (0.75 in.) supply hoses, with a 300 kW CDU.

Use shutoff valves for each supply line that exits the manifold to enable the flow of water to be stopped in individual legs of multiple-circuit loops. The shutoff valves provide a way of servicing or replacing an individual heat exchanger without affecting the operation of other heat exchangers in the loop.

Use adjustable flow control valves, called circuit setters, on each supply line to enable control of the flow and compliance with specifications for each heat exchanger.

Use temperature and flow metering (monitoring) in secondary loops to provide assurance that water specifications are being met and that the optimum heat removal is taking place.

Use circuit setters, placed as shown in Figure 5 and Figure 6, to enable the adjustment of water flow to each water circuit off a main manifold system.

Anchor or restrain all manifolds and pipes to provide the required support and to avoid movement when quick-connect couplings are being attached to the manifolds.

Figure 5 shows a layout for multiple water circuits.

Figure 5. Typical central manifold (at a central location for multiple water circuits)

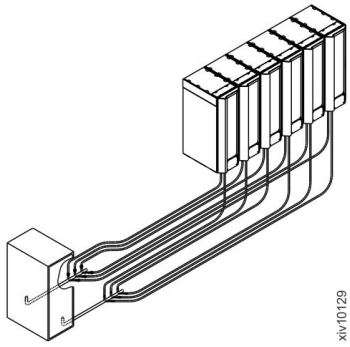

Figure 6 shows an extended manifold layout.

Figure 6. Typical extended manifold (located along aisles between racks)

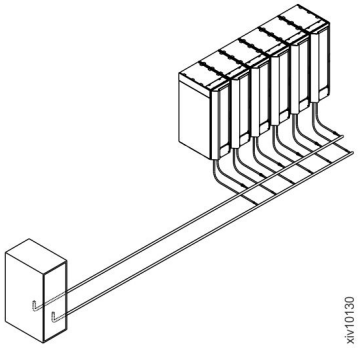

Flexible hoses and connections to manifolds and heat exchangers

Pipe and hose configurations can vary. You can determine the best configuration for your installation by analyzing the needs of your facilities, or a site preparation representative can provide this analysis.

Flexible hoses are needed to supply and return water between your hard plumbing (manifolds and cooling distribution units) and the heat exchanger, allowing needed movement for opening and closing the rack rear door.

Hoses are available that provide water with acceptable pressure-drop characteristics and that help prevent depletion of some corrosion inhibitors. These hoses must be made of peroxide cured ethylene propylene diene monomer (EPDM), non-metal-oxide material and must have Parker Fluid Connectors quick-connect couplings at each end. These couplings are compatible with the heat exchanger couplings. Hose lengths from 3 to 15 meters (10 - 50 ft), in increments of 3 meters (10 ft), are available. Hoses that are longer than 15 meters (50 ft) might create unacceptable pressure loss in the secondary circuit and reduce the water flow, reducing the heat removal capabilities of the heat exchanger.

For information about a supplier of these hoses, see Secondary cooling loop parts and services information. Use solid piping or tubing that has a minimum inner diameter of 19 mm (0.75 in.) and the fewest possible joints between a manifold and a heat exchanger in each secondary loop.

You can use quick-connect couplings to attach the hoses to the distribution manifolds. You must use quick-connect couplings to attach the hoses to the heat exchanger. Hose couplings that connect to the heat exchanger must have the following characteristics:

- The couplings must be constructed of passivated 300-L series stainless steel or brass with less than 30% zinc content. The coupling size is 19 mm (0.75 in.).
- The supply hose must have a Parker (male) quick-coupling nipple, part number SH6-63-W, or equivalent. The return hose must have a Parker (female) quick-connect coupling, part number SH6-62-W, or equivalent.
- At the opposite (manifold) end of the hoses, use similar quick-connect couplings. However, if you want to use other types of connectors, make sure that positive locking mechanisms are used to prevent loss of water when the hoses are disconnected. The connections must minimize water spill and air inclusion into the system when they are disconnected. You can also permanently attach hoses to the manifolds by using barbed fittings and clamps.

Raised floor hose requirements and management

On a raised floor, hoses can be routed under the floor tiles and can be brought up from beneath the rack through special tile cutouts. The hoses are connected to the quick-connect couplings on the bottom of the heat exchanger.

In a typical example, each heat exchanger requires a special cut 0.6 m × 0.6 m (2 ft × 2 ft) floor tile below it with the opening outside of the rack footprint. A portion of the tile is cut away and correctly covered to protect against sharp edges. The corner opening is placed directly under the hinge side of the rack rear door. The opening size of the cut is 152.4 mm wide × 190.5 mm long ±12.7 mm (6.0 in. wide × 7.5 in. long ±0.5 in.) in the direction parallel to the door. See Figure 1 and Figure 2.

Figure 1. Raised-floor hose management example 1: hose exit through floor tile at the door hinge

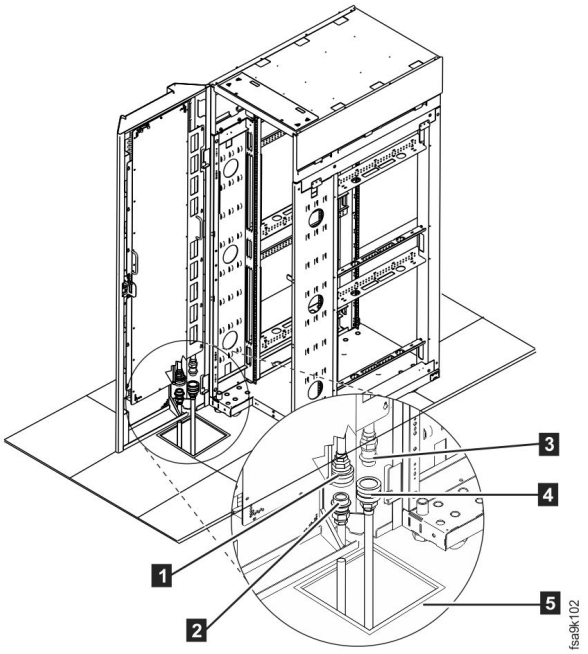

- 1 Heat exchanger (female coupling)
- 2 Supply hose assembly (male)
- 3 Heat exchanger (male coupling)
- 4 Return hose assembly (female)
- 5 Raised floor

Figure 2. Raised-floor hose management example 2: tile cutout size and position

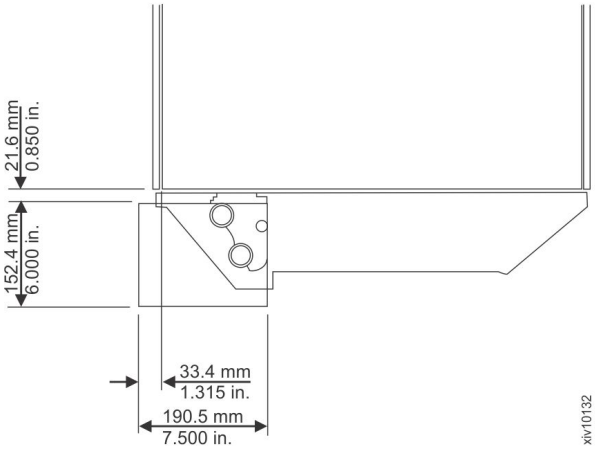

In another example, for a rack that is installed at the same time as a heat exchanger, or in cases where a rack is moved to install new floor tiles under it, each heat exchanger still requires a special cut 0.6 m × 0.6 m (2 ft × 2 ft) floor tile. However, the floor tile is positioned completely within the footprint of the rack. A modified cable opening or independent hose cutout is used.

Flexible hoses that each contain a right-angle elbow are used to route the hoses under the rack in a large loop to allow hose movement when the door is opened and closed. Figure 3 and Figure 4 show how to route hoses under the rack with enough hose length to allow the hose to move freely as the door is opened and closed.
Note: Existing tile cutouts for electrical or other cables can also be used for the hoses, if enough space is available.

Figure 3. Raised-floor and non-raised-floor hose management example 2:loop under the rack with door closed

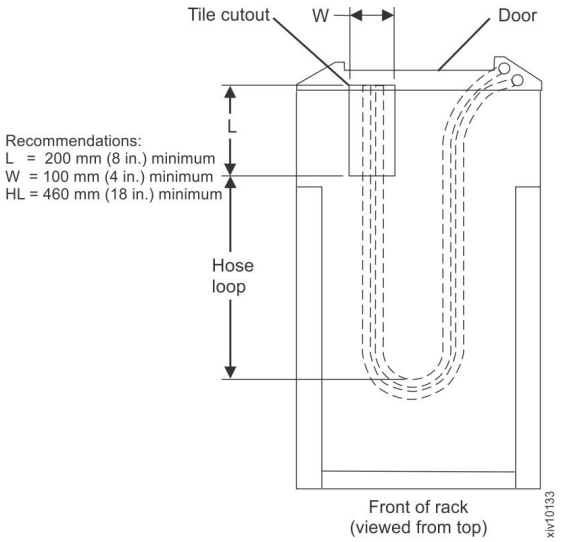

Figure 4. Raised floor and non-raised floor hose management example 2:loop under the rack with door open

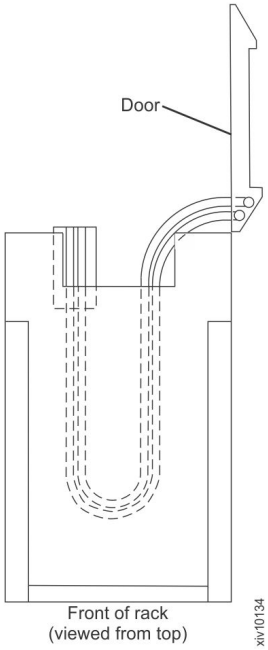

Lay hoses side-by-side as they run between the heat exchanger and the pump unit manifold, and allow the hoses to freely move. Leave enough slack in the hoses below the rear door so that no pressure is exerted on the mated couplings when the hoses are connected and operating. When you route hoses, avoid sharp bends that cause hose kinks and avoid hose contact with sharp edges.

Non-raised floor hose requirements and management

In data centers without a raised floor, straight hose assemblies cannot make the sharp bend to exit between the floor and the rack door without kinking the hose.

Hose assemblies with right-angle metal elbows are needed to route the hoses along the floor. Make the 90° turn upwards within the gap between the bottom of the heat exchanger and the floor surface, and then connect to the heat exchanger couplings (see Figure 1).

Figure 1. Non-raised floor hose requirements

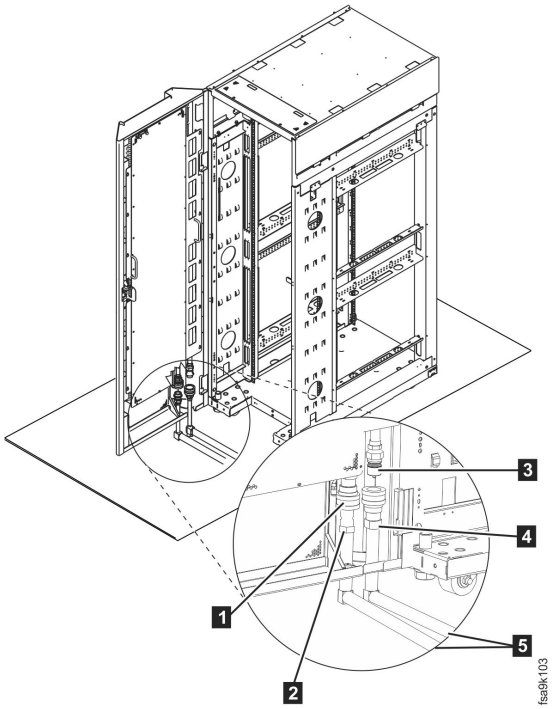

- 1 Heat exchanger (female coupling)
- 2 Supply hose assembly (male)
- 3 Heat exchanger (male coupling)
- 4 Return hose assembly (female)
- 5 Elbow extensions

Hoses that exit the heat exchanger are routed in a manner similar to that of power cables in a non-raised-floor data center. For example, place the hoses side-by-side and allow them to move freely as they approach the rack (within approximately 3 meters [10 feet] of the rack).

When you open the door, it is acceptable for the hoses to move slightly and rotate in parallel at the coupling interface inside the door. As you close the door, the hoses rotate back to their original positions.
Note: When you open or close the door, some manipulation of the hose along the floor might be necessary to prevent unwanted forces on the door and to make it easier to open and close the door.
Hose coverings or protective devices are not provided by IBM®. Routing and protection of the hose assemblies exterior to the rack are your responsibility.

Secondary cooling loop parts and services information

IBM® supplies a heat exchanger that is designed for IBM enterprise-server racks, a hinge kit (for those racks), and an air purge tool. This information provides sources and information for other parts and services that are needed for correct function and reliability of the secondary cooling loop.

Servicing and miscellaneous parts supplier

Table 1 provides supplier and contact information for miscellaneous secondary loop parts. You can contact the supplier that is listed in the table for all or some of the items that are listed, depending on your needs.

Table 1. Servicing and miscellaneous secondary loop parts supplier information for customers in North America, Europe, Middle East, Africa, Asia Pacific

Supplier	Parts	Contact information
Wakefield-Vette, Inc.	• Rear door heat exchangers (designed for both non-IBM as well as IBM Enterprise racks) • Cooling distribution units • Hose kits • Water treatment • Chillers • Raised-floor grommets	**Website**: Wakefield-Vette, Inc. **Phone**: 603- 635-280-0877 **Corporate Headquarters**: Wakefield-Vette 33 Bridge Street Pelham, NH 03076

Services supplier

Table 2 provides supplier and contact information for services that can be provided for secondary loop parts.

Table 2. Services supplier information for customers in North America, Europe, Middle East, Africa, Asia Pacific

Supplier	Services	Contact information
Wakefield-Vette, Inc.	• Installation of door and secondary loop items • Preventive maintenance	**Website**: Wakefield-Vette, Inc. **Phone**: 603- 635-280-0877 **Corporate Headquarters**: Wakefield-Vette 33 Bridge Street Pelham, NH 03076

Cooling distribution unit suppliers

Table 3 provides supplier and contact information for customers in Europe for a cooling distribution unit that was designed specifically for the IBM Rear Door Heat exchanger.
Note: Customers in other locations can contact Nortek Global HVAC (formerly Eaton-Williams Group) or Wakefield-Vette (see Table 2).

Table 3. Cooling distribution unit supplier information for customers in Europe

Supplier	Cooling distribution units	Contact information
Nortek Global HVAC (formerly Eaton-Williams Group, Ltd (UK))	• CDU120 (120 kW, 400 - 480V) • CDU121 (120 kW, 208 V) • CDU150 (150 kW, 400 - 480V) • CDU151 (150 kW, 208 V)	**Website**: Nortek Global HVAC

The following illustration shows the Eaton-Williams cooling distribution unit features.

Figure 1. Eaton-Williams cooling distribution unit features

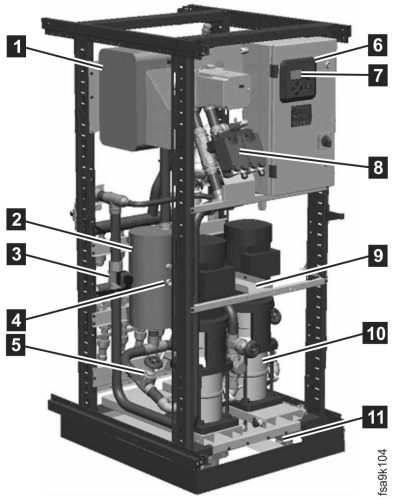

- 1 Plate heat exchanger
- 2 Reservoir tank
- 3 Primary flow meter
- 4 Level switch
- 5 Pump isolation valves
- 6 Control panel
- 7 Control display (user interface)
- 8 Primary control valves (run and standby)
- 9 Pump transit bracket (remove after installation)
- 10 Secondary pumps (run and standby)
- 11 Lockable castor and adjustable plinth

The cooling distribution unit specifications are described in Table 4.

Table 4. Eaton-Williams cooling distribution unit specifications

Performance	Properties
Maximum cooling capacity	120 kW (409 450 Btu/hr) *or* 150 kW (511 815 Btu/hr)
Pump capacity (design flow)	240 L/min (63.4 GPM)
Maximum pump head pressure	355 kPa (51.5 psi) at design duty, excluding cabinet losses
Coolant (liquid) type	Chilled water (with up to 30% glycol) Important: Primary facilities input only. Note recommended for secondary loop, as noted in Water control and conditioning for the secondary cooling loop.
Primary liquid connections	1 1/2 in. flex tail for sweat connection, top or bottom

Performance	Properties
Secondary liquid connections	3/4 in. quick connects, hydraulic ISO-B
Unit internal primary circuit liquid capacity	Approximately 10.0 liters (2.6 gallons)
Unit internal secondary circuit liquid capacity	Approximately. 32.0 liters (8.5 gallons)
Noise	< 55 dBA at 3 meters
Electrical	
Power supply	200 - 230 V, 3Ø, 50/60 Hz *or* 400® - 480 V, 3Ø, 50/60 Hz
Maximum power consumption	5.6 kVA at 480 V, 4.9 kVA at 208 V
Physical	
Height	1825 mm (72 in.)
Width	800 mm (31.5 in.)
Depth	1085 mm (43 in.)
Weight (empty)	396 kg (870 lb)
Weight (filled)	438 kg (965 lb)

Note: Other industrial cooling distribution units can be used in a secondary cooling loop with the rear-door heat exchanger, if they meet the specifications and requirements that are described or referred to in this document.

Maintenance schedule

Maintenance tasks are required at certain time intervals.

Task	Schedule
Check the manifolds for temperature (make sure that the top of the manifolds is cool) and sounds of air in the system to ensure that the exchanger is correctly filled.	One month after installation and again annually
Inspect the heat-exchanger fins for air blockage at the fins (such as dust, dirt, and debris)	Annually
Inspect the entire length of the supply hose and return hose for damage, age cracks, and kinks. Be sure to inspect at the door and outside of the rack.	Annually

Site security considerations

When installing an IBM FlashSystem® A9000R storage system, you need to apply the same security practices that you apply to any other business critical IT system.

Note: A good reference on storage security can be found at the Storage Networking Industry Association (SNIA) website: http://www.snia.org/forums/ssif/programs/best_practices.
A common risk with storage systems is the retention of volatile caches. Your FlashSystem A9000R storage system is perfectly safe in regard to external operations and a loss of external power. If there is a power failure, the internal battery backup modules provide power to the system. These battery modules allow your storage system to gracefully shut down.
If desired, you can install your own uninterruptible power supply (UPS) unit(s) or generators in order to provide further power-failure protection.

However, if someone gains physical access to the equipment, physical contact with the machine presents the following security risks:

- Stealing of the machine or components
- Breaking components
- Manual shutdown of the machine or components by bypassing the preferred process

These cases could lead to losing the contents of the system and its volatile caches, resulting in loss of access.
Attention: Restricting physical access is especially critical when using local key management.
For more information on using local key management, see Native user authentication.

For more information on using external key management (LDAP), see External authentication via Lightweight Directory Access Protocol (LDAP).

To eliminate or greatly reduce this risk, the IBM FlashSystem A9000R rack can be equipped with lockable doors (feature code AFR6).
Important: In addition to lockable doors on the system itself, it is highly advised to secure access to the room in which the physical storage system is kept.

Network and host connectivity requirements

This section details the host SAN connections and requirements for IBM FlashSystem® A9000R storage systems.

- **Network connections for management**
  Network connectivity for the storage system is provided through the utility patch panel. Use this information to identify patch panel connectivity.
- **Management port requirements**
  The management ports provide the connectivity required for the IBM® Hyper-Scale Manager, the IBM FlashSystem A9000R command-line interface (CLI), and other management tools to monitor and control the FlashSystem A9000R system.

- **Host network connectivity and multipathing configurations**
  Host systems can connect to an IBM FlashSystem A9000R over a Fibre Channel (FC) network by using the Small Computer System Interface (SCSI) protocol or over an Ethernet network by using the Internet Small Computer System Interface (iSCSI) protocol.
- **Network cable requirements**
  The customer is responsible for supplying cables that connect to the patch panel, including host attachment (Fibre Channel or iSCSI) cables, management cables, maintenance cables, and virtual private network (VPN) cables.
- **Network and host connectivity security information**
  The storage system integrates various security features, to protect your network and host systems.

Network connections for management

Network connectivity for the storage system is provided through the utility patch panel. Use this information to identify patch panel connectivity.

- Figure 1 illustrates the utility patch panel for model 415.
- Figure 2 illustrates the utility patch panel for models 425 and U25.
- Table 1 describes the ports that are available in the patch panel and the component to which each patch panel port connects.

Figure 1. Utility patch panel – model 415

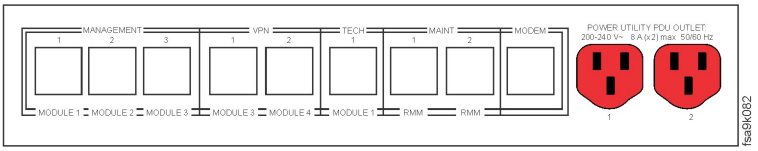

Figure 2. Utility patch panel – models 425 and U25

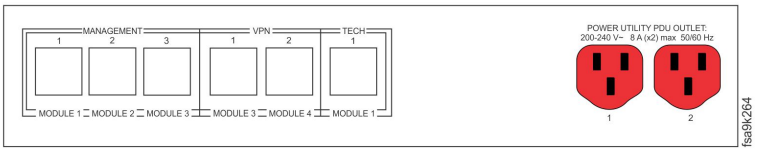

Table 1. Utility patch panel connections

Patch panel port	Component	Description
Management ports	Grid controllers 1, 2, 3	Use these three redundant ports to connect to systems that are used for managing the IBM FlashSystem A9000R system using the IBM® Hyper-Scale Manager UI (UI) and command-line interface (CLI). These ports can also be used for sending email notifications and SNMP traps about event alerts, NTP server time synchronization, and communication with key servers for managing encryption keys. Important: Use all three connections for redundancy. If you use only one path and the grid controller for that connection goes down, the ability to manage the IBM FlashSystem A9000R system is lost. For information about IP configuration and protocol requirements, see Management port requirements.
VPN (remote support) ports	Grid controllers 3, 4	These two virtual private network (VPN) ports are used to connect to the IBM FlashSystem Remote Support Center. Note: • Use both VPN connections for redundancy to avoid losing remote support if one of the grid controllers fails. • Some model 425 or U25 minimal rack configurations may only contain three grid controllers. In this case, only use grid controller 3 VPN connectivity.
Tech (technician) ports	Grid controller 1	(For technician use only.) This port is used by the IBM service representative to connect a notebook computer to the storage system, for initial configuration and servicing the system. A Dynamic Host Configuration Protocol (DHCP) server is implemented over this port. The DHCP server automatically assigns IP addresses to the notebook computer and establishes a connection to the IBM FlashSystem A9000R system. Note: There is a second identical technician port on the front of the rack, next to the power button. This technician port connects to grid controller 2.
Maint (maintenance module) ports (Model 415 only)	Maintenance module	(For technician use only.) Not in use.
Modem port (Model 415 only)	Modem	(For technician use only.) Not in use.

Management port requirements

The management ports provide the connectivity required for the IBM® Hyper-Scale Manager, the IBM FlashSystem® A9000R command-line interface (CLI), and other management tools to monitor and control the FlashSystem A9000R system.

- **IP configuration**
  Use this information for system storage IP configuration.

- **Control (management) protocols**
  Communication to the system via Ethernet ports is divided into two: data Ethernet ports and control Ethernet ports. The control protocols are used for system management; monitoring; telemetry and support; and communication security. These control protocols are detailed in this section.

IP configuration

Use this information for system storage IP configuration.

The IBM FlashSystem® A9000R storage system has three redundant management port IP addresses, over different Ethernet interfaces, in case of failure.

Each IP address is handled by different grid controllers in the initial grid element configuration.
Note: If you use only one path and the grid controller for that connection goes down, the ability to manage the system is lost.
For information about the IBM® Hyper-Scale Manager, see the IBM Hyper-Scale Manager documentation .

Management functions can be performed through any of the IP addresses. These addresses can be accessed simultaneously by multiple clients.

For each of the three management ports, provide the following information to the IBM service representative before starting the installation process:

- IP address of the port
- Net mask
- Default gateway IP
- Maximum transmission unit (MTU)
  Attention: Be careful when changing an Ethernet port MTU. If management connectivity is on, or a VPN is being used, this might break existing application connectivity unless the adjacent switch is configured properly.
  This would be the result of the switch using a fixed MTU size lower than the new size being defined on the system.

  For example, if the system sends 9000 byte packets while the switch can only receive packets up to 1500 bytes, the switch drops extra packets.

  Note:
  - All management IP interfaces must be connected to the same subnet and use the same network mask, gateway, and MTU.
  - MTU configuration is required if the network supports an MTU that is greater than the default 1536 bytes. The largest possible MTU supported is 9216 bytes.
- Virtual LAN (VLAN)
  VLAN tagging can be enabled or disabled for management and VPN ports. The ports can be configured as untagged (default) or tagged.
  Note:
  - A valid VLAN tag value can only be 1 - 4094.
  - All management or VPN ports with VLAN tagging enabled must belong to the same VLAN tag, one per port type.

In addition, provide the following system-level IP information:
Note: IPV6 is supported.

- IP address of the primary and secondary DNS servers
- IP addresses or DNS names of the SMTP servers
- IP addresses or DNS names of the NTP server
- IP addresses and port numbers of the Remote Support servers

Control (management) protocols

Communication to the system via Ethernet ports is divided into two: data Ethernet ports and control Ethernet ports. The control protocols are used for system management; monitoring; telemetry and support; and communication security. These control protocols are detailed in this section.

Note: For more information regarding the data Ethernet configurations, see iSCSI network configurations.
System management is carried out through the protocols detailed in Table 1.
Note: *Inbound* protocol connections are initiated by external entities, such as IBM Hyper-Scale Manager and others. The *outbound* protocol connections are initiated by FlashSystem A9000R.

Table 1. Control (management) protocols

Use	Protocol	TCP port	Inbound/ Outbound	Encrypted?	Enabled by default (Yes/No)	Comments
Management						
Command-line interface (CLI) (includes Hyper-Scale Manager)	<proprietary>	7778	Inbound	Yes	Yes	Storage system management portal and CLI act as the client and initiate the connection, while FlashSystem A9000R acts as the server.
Common Information Model (CIM) agent	CIM	5989	Inbound	Yes	No	Storage system uses CIM agent to reply to management commands.
Common Information Model (CIM)	SLP	427	Inbound	No	No	Service Location Protocol (SLP) used for service discovery in conjunction with the CIM protocol.
Time setting	NTP	123	Outbound	No	N/A	Storage system uses a Network Time Protocol (NTP) connection.

Use	Protocol	TCP port	Inbound/ Outbound	Encrypted?	Enabled by default (Yes/No)	Comments
Monitoring						
SNMP requests	SNMP	161	Inbound	Yes	Yes	Storage system responds to SNMP requests when sending replies to SNMP managers.
SNMP traps/notifications	SNMP	162	Outbound	No	N/A	Storage system initiates SNMP messages when sending traps/notifications to SNMP managers.
Syslog	UDP	514	Outbound	No	N/A	Storage system uses for collecting system logs.
Telemetry and Support						
Outbound mail	SMTP	25	Outbound	No	N/A	Storage system initiates SMTP traffic when sending emails for either event notifications or for SMS gateways. Note: Closing this port disconnects the IBM Service Center.
XRSC (SSH)	TCP	22, 2222	Inbound	Need to enable	Yes	Storage system uses port for remote connectivity. Note: The management ports should be on a different subnet than the VPN ports that are used for remote access.
Secure technician access	TCP	443	Inbound	Yes	No	Technician Assistant tool uses for data management platform (DMP) connectivity.
Communication security and access						
Connectivity	IPSec	1293	Inbound	Yes	No	Storage system uses IPSec for management and VPN communication. Note: The management ports should be on a different subnet than the VPN ports that are used for remote access.
Domain Name System	DNS	53	Outbound	No	N/A	Storage system uses a Domain Name System (DNS) connection.
LDAP authentication	LDAP	389	Outbound	No	N/A	Storage system uses for LDAP authentication connectivity.
Other						
Key management	KMIP	5696	Outbound	Yes	N/A	Storage system communicates with key servers using the KMIP protocol.
Quorum Witness	<proprietary>	8460	Outbound	Yes	N/A	Storage system uses port with Quorum Witness for API communication (with a default Linux firewall).
Quorum Witness	<proprietary>	8461	Outbound	Yes	N/A	Storage system uses port with Quorum Witness for log retrieval (with a default Linux firewall).
Quorum Witness	<proprietary>	8462	Outbound	Yes	N/A	Storage system uses port with Quorum Witness for retrieval of Ngnix statistics (with a default Linux firewall).

Host network connectivity and multipathing configurations

Host systems can connect to an IBM FlashSystem® A9000R over a Fibre Channel (FC) network by using the Small Computer System Interface (SCSI) protocol or over an Ethernet network by using the Internet Small Computer System Interface (iSCSI) protocol.

Important: A host must be attached to FlashSystem A9000R system through a Fibre Channel fabric or Ethernet switch.
While a host can connect through FC and iSCSI simultaneously, the same LUN can only be mapped through FC *or* iSCSI.
Host traffic can be directed to any of the grid controllers.

The administrator must ensure that multipathing is used. Multipathing is configured by ensuring the following:

- Host connections avoid single points of failure by applying redundant connections.
- All host workload is adequately balanced across the connections and grid controllers (ensuring system resource utilization is maximized). Review the balancing periodically and when connections or traffic patterns change.

Important: Always employ multipathing for data connections. If multipathing is not used, the Health widget on the Dashboard of the UI indicates that multipathing is not used.
For more information about host attachment support and connectivity, see Configuring host attachment. For supported interoperability configurations, see the System Storage® Interoperation Center website .

- **Fibre Channel (FC) network configurations**
  Host systems can connect to the storage system over a Fibre Channel (FC) network. Use these important practices when configuring your Fibre Channel host port connections, in order to achieve high availability and high performance in your storage system.
- **iSCSI network configurations**
  Host systems can connect to a FlashSystem A9000R over an Ethernet network using the Internet Small Computer System Interface (iSCSI) protocol. Use these important recommendations when configuring your iSCSI host port connections, in order to achieve high availability and high performance in your storage system.

Fibre Channel (FC) network configurations

Host systems can connect to the storage system over a Fibre Channel (FC) network. Use these important practices when configuring your Fibre Channel host port connections, in order to achieve high availability and high performance in your storage system.

- **Fibre Channel (FC) adapter types**
  Use this information to learn about the FC adapter types being used in the IBM FlashSystem A9000R systems.
- **Fibre Channel host port configuration**
  Use this information to properly configure your Fibre Channel host port connectivity.
- **Fibre Channel connectivity requirements**
  Use these connectivity requirements in order to ensure redundancy, further protecting your system data.
- **Fibre Channel best practices**
  Use this information for FC host connectivity best practices.
- **Fibre Channel zone types**
  A *zone type* is defined as a set of targets from a certain combination of grid controllers. Each zone defined in each fabric is based on a single *zone type*.
- **Fibre Channel connectivity during scale-out (MES)**
  Use this information for preparing host connectivity during system scale-out (MES).

Fibre Channel (FC) adapter types

Use this information to learn about the FC adapter types being used in the IBM FlashSystem® A9000R systems.

Both NVMe/FC and FC-FCP adapter types offer the same 16 Gb bandwidth.

The external physical difference between the two adapers is as follows:

- FC-FCP adapters are stamped with "PCIe x8 16GbFC".
  The FC-FCP port markings are stamped on each adapter.
- NVMe/FC adapters are stamped with "PCIe FC".
  The NVMe/FC port numbering is labeled on each adapter.

Figure 1 shows the placement of the FC markings.

Figure 1. NVMe/FC adapter indication

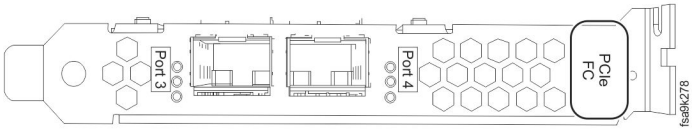

Note: Figure 1 shows an example of an NVMe/FC adapter in slot 8. The adapters in slot 4 looks slightly different, but all of the visual indicators described and shown in the image are present.
To determine which adapters are installed in your system, lookup the grid controller feature code, using the feature code information in Grid controller components and feature codes.

NVMe/FC adapters

NVM Express (NVMe) allows servers to leverage the native parallelism of today's SSD offerings, reduces overall I/O overhead, and increases bandwidth. NVMe/FC enables NVMe over a Fibre Channel (FC) network fabric, thus combining the benefits of all-flash SAN storage, with NVMe performance, over existing infrastructure.

Starting from system code level 12.2.1, all new orders of IBM FlashSystem A9000R models 425 and U25 systems (system code level 12.2.1 and higher) are shipped with an enhanced grid controller, with NVMe/FC ready adapters.

From system code level 12.4.0, all systems containing grid controller feature code 5005 have the ability to connect to either NVMe-based or SCSI-based (FCP) switches.

For systems with software versions prior to 12.4.0 and purchased prior to 12.2.1 release, the NVMe/FC adapters serve FC over SCSI only, just as the FC-FCP adapters currently do. In this type of FC configuration, a *system* may contain FC-FCP adapters; NVMe/FC adapters; or both. However, a *single grid controller* may contain only *one type* of adapter.

FC-FCP adapters

These adapters are used on grid controllers for model 415 and any model 425 that was proposed or ordered prior to 12.2.1.

Fibre Channel host port configuration

Use this information to properly configure your Fibre Channel host port connectivity.

Physical port connectivity

Be sure to connect port 1 of each grid controller in the system to switch 1 which belongs to fabric 1, and port 3 of each grid controller to switch 2 which belongs to fabric 2.
For more connectivity information regarding port connectivity, see Fibre Channel connectivity requirements. For a port–fabric connectivity example, see Figure 2.

Connectivity for high performance

For high performance, follow these important practices:

- Spread grid controller port connections as evenly as possible.
- For high CPU utilization in each grid controller, use ports 1 and 3 for host connectivity and ports 2 and 4 for migration and mirroring activities (see Migration and mirroring connectivity).

Ensuring high availability of each host

In order to ensure high availability of each host, follow these important practices:

- Ensure that you have more than one path from the host to the system.
- Divide the paths between the two fabrics: half the paths connecting to fabric 1 and the other half connecting to fabric 2.
- Spread the paths across different grid controllers.
- Use different adapters within each of the grid controllers, as shown in Figure 1.

Figure 1. Crossing Fibre Channel grid controller ports

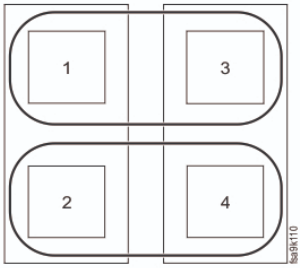

Important: It is important to use both Fibre Channel adapters for resiliency and high performance.
Figure 2 shows the port numbers for each of the FC ports on the rear of the grid controllers.

Figure 2. Fibre Channel port numbering on the grid controllers

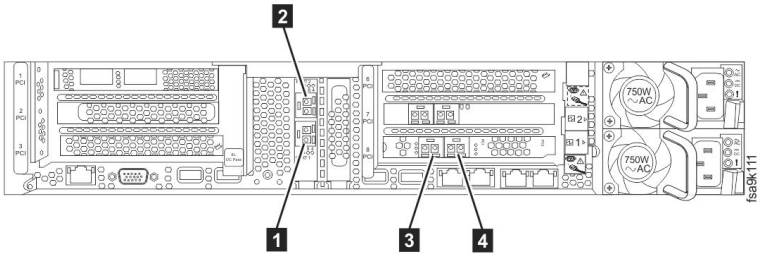

- 1 FC port 1
- 2 FC port 2
- 3 FC port 3
- 4 FC port 4

Only use non-redundant configurations when the risks of a single point of failure are acceptable, which is typically the case for test and development environments. Non-redundant configurations should generally not be used. For connectivity requirements, see Fibre Channel connectivity requirements.

Fibre Channel connectivity requirements

Use these connectivity requirements in order to ensure redundancy, further protecting your system data.

In a production environment, always connect Fibre Channel hosts to a **minimum** of two independent fabrics.
For rack configurations with two or more grid elements, be sure to minimally have four separate grid controller connections, two paths to each fabric, as illustrated in Figure 1.

Figure 1. Minimum host connectivity for two or more grid elements

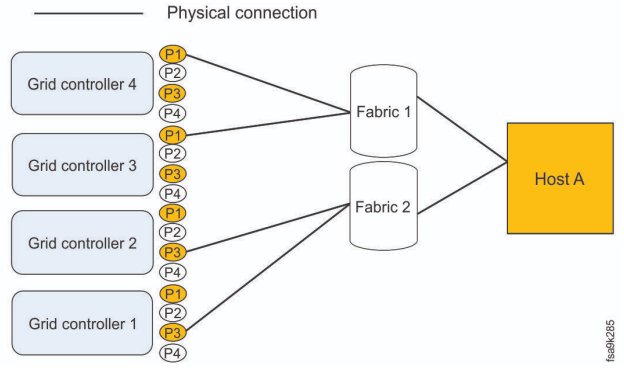

For model 425 or U25 minimum racks with three grid controllers, be sure to minimally have two separate grid controller connections, each path going to a separate fabric, as illustrated in Figure 2.

Figure 2. Minimum host connectivity for model 425 and U25 minimum racks

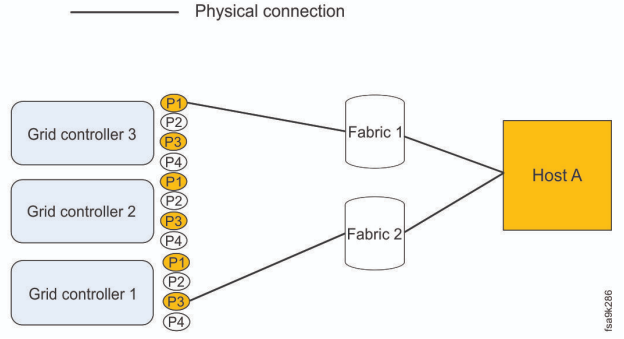

Important: Host multipath connectivity eliminates the risk of a single point-of-failure between the host and storage systems.
In Fibre Channel network configurations, there are two dual-port connections per grid controller. This setup contains a total of four paths (two paths to each fabric), providing protection from fabric failure and failure of up to two grid controllers.
For models 425 and U25 minimum configuration containing three grid controllers, the setup contains a total of two paths (one path to each fabric), providing protection against single fabric failure or a failure of up to one grid controller.
Note:

- For best performance, use a 16 Gb Fibre fabric and HBA on the host.
- Host system can have as many HBAs as needed.
- In a FlashSystem A9000R with Fibre Channel configuration, each grid controller has two dual 16 Gb Fibre Channel ports.

Several network configurations that use Fibre Channel are technically possible, and each configuration varies in terms of cost, flexibility, performance, and reliability.

Fibre Channel best practices

Use this information for FC host connectivity best practices.

Talk to your IBM® representative to help optimally setup your host connections, according to your needs.

To achieve high performance, it is important to:

1. Spread the host connections to each grid controller evenly.
2. Utilize the CPU in each grid controller as much as possible. Use ports 1 and 3 for host connectivity and then ports 2 and 4 for migration and mirroring.

To achieve balancing, all hosts should be spread across all *zone types* evenly. The amount of *zone types* depends on the number of grid controllers in the system being used for connectivity.

For more information on *zone types*, see Fibre Channel zone types.

For supported interoperability configurations, see the System Storage® Interoperation Center website .

Fibre Channel zone types

A *zone type* is defined as a set of targets from a certain combination of grid controllers. Each zone defined in each fabric is based on a single *zone type*.

Zoning that follows the *zone type* configuration should guarantee the following:

- High availability of host connectivity.
- All grid controllers and ports are used evenly (as described in Fibre Channel host port configuration).
- Once the *zone type* is defined, it should not be changed on any scale-out (MES) situations.
  For information on host connectivity during scale-out, see Fibre Channel connectivity during scale-out (MES).

While Host A may be associated to Zone 1 and Host B may be associated to Zone 2, they are physically connected to the same system ports on different zones, therefore they have the same *zone type*. Zone type for 2 grid elements (4 grid controllers) shows an example of this type of zoning scheme.
For more information on *zone types*, zoning, connectivity during scale-out, and general host connectivity, see IBM® FlashSystem A9000, IBM FlashSystem® A9000R, and IBM XIV® Storage System: Host Attachment and Interoperability Redbook, SG24-8368 .

- **Zone type examples**
  Use the following *zone type* examples to properly configure your Fibre Channel host connectivity.

Zone type examples

Use the following *zone type* examples to properly configure your Fibre Channel host connectivity.

Zone type for model 425 and U25 minimal rack (3 grid controllers)

A model 425 or U25 minimal rack with three grid controllers will have only one (1) *zone type* (Z123). Physically each grid controller is connected to two different fabrics.

Table 1. Map of ports according to *zone type* – models 425 and U25 minimal rack

Zone type name	Grid controller 1	Grid controller 2	Grid controller 3
	Ports 1 and 3	Ports 1 and 3	Ports 1 and 3
Z123	✓		

Figure 1. Single *zone type* for model 425 and U25 minimum racks

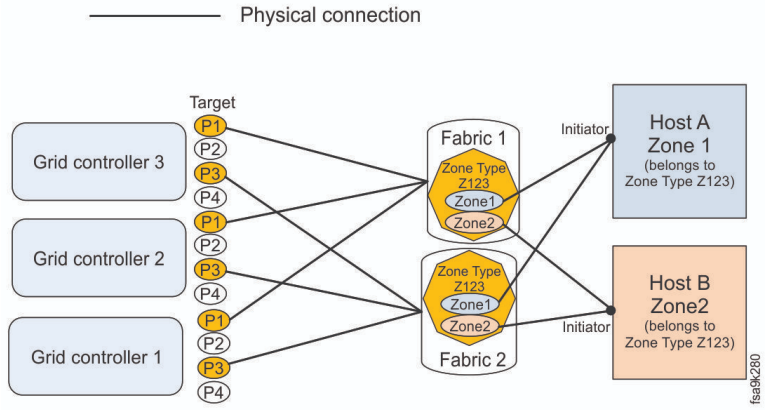

Zone type for 2 grid elements (4 grid controllers)

A two grid element (minimal rack for model 415) will have only one (1) *zone type* (Z1234). Each host is attached to the system through a zone which is defined by this *zone type* (Z1234).
Physically each grid controller is connected to two different fabrics through a single zone, being defined on both of the fabrics. This particular zone is based on a single *zone type*.

Table 2 shows how all ports across all grid controllers are configured within the same *zone type*. Figure 2 illustrates the port connections across the different hosts and fabrics, with zoning examples, where P1 and P3 are the port number connections.

Table 2. Map of ports according to *zone type* – 2 grid elements

Zone type name	Grid controller 1	Grid controller 2	Grid controller 3	Grid controller 4
	Ports 1 and 3	Ports 1 and 3	Ports 1 and 3	Ports 1 and 3
Z1234	✓			

Figure 2. Single *zone type* for 2 grid elements

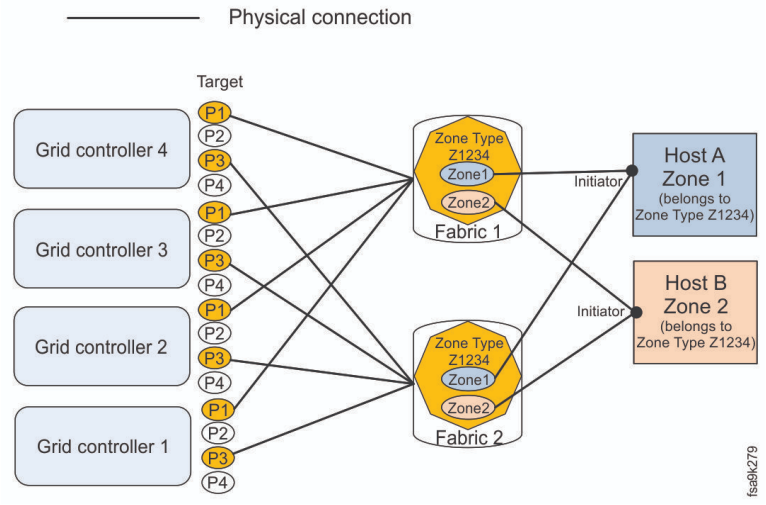

Zone type for 3 grid elements (6 grid controllers)

When a system has six (6) grid controllers, three (3) *zone types* should be used. Host zoning is spread across the *zone types* evenly, in order to achieve balancing. Each host is based on a specific *zone type* in such a way that all host zones are spread across all the *zone types*.

Table 3 shows how the ports on the different grid controllers are configured in three different *zone types*. Figure 3 illustrates the port connections across the different hosts and fabrics, according to *zone type*, where P1 and P3 are the port number connections.

Table 3. Map of ports according to *zone type* – 6 grid controllers

Zone type name	Grid controller 1	Grid controller 2	Grid controller 3	Grid controller 4	Grid controller 5	Grid controller 6
	Ports 1 and 3	Ports 1 and 3	Ports 1 and 3	Ports 1 and 3	Ports 1 and 3	Ports 1 and 3
Z1234	✓					
Z3456			✓			
Z1256	✓				✓	

Figure 3. Zoning example for 3 grid elements

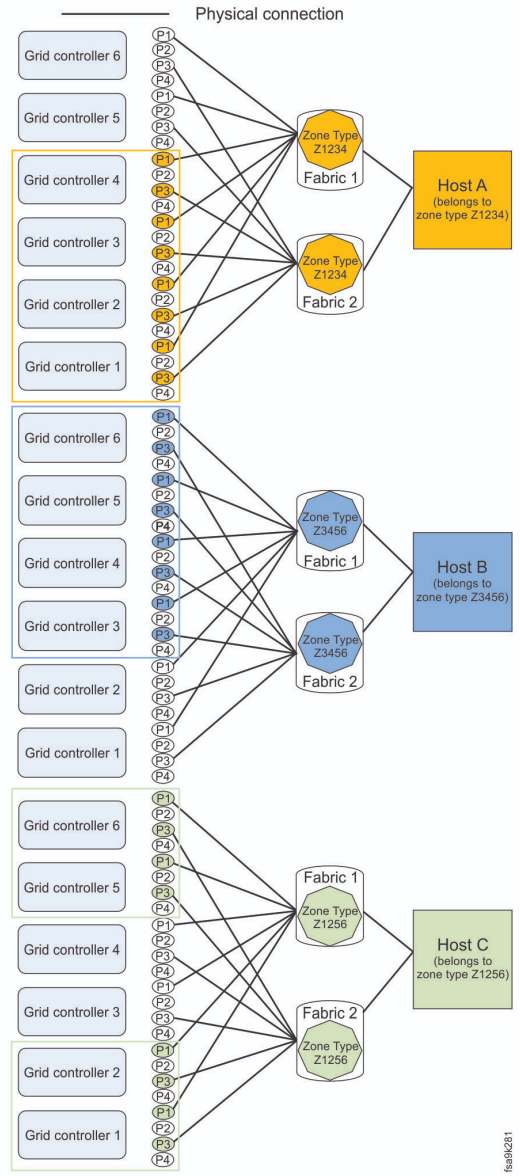

Fibre Channel connectivity during scale-out (MES)

Use this information for preparing host connectivity during system scale-out (MES).

For more information about system scale-out, see Conducting system scale-out (MES).

- The original physical port connections are kept.
- Continue the physical port connection divisions, as described in Fibre Channel host port configuration.
- *Zone types* might need to be added or removed, depending on the number of grid controllers that were originally connected to the system and how many are being added during scale-out.
  Important: In order to preserve proper balancing, be sure that the host zones remain spread evenly across all *zone types* when reconfiguring zoning during scale-out.
  Note: During scale-out, in order to maintain a proper connectivity balance, it may be necessary to remap some of the hosts that were previously within one *zone type* to be mapped into one of the new *zone types*.
- Be sure to add the new targets before deleting the ones that are no longer necessary.
  Use the following steps for reconfiguring the zoning:
  1. Add the new targets.
  2. Rescan the ports in the host.
  3. Delete the targets from the host side that are no longer needed.
  4. Rescan the host ports.
  5. Verify that all paths are working properly.

- **Example of adding zone types during scale-out**
  Use this example as a guideline of adding *zone types* during scale-out.

Example of adding zone types during scale-out

Use this example as a guideline of adding *zone types* during scale-out.

The following is an example of adding *zone types* during a scale-out scenario of adding one (1) grid element to an existing system with two (2) grid elements (a minimal system).

In this case, there was a single zone type (Z1234) to which all of the existing hosts were mapped. Now, some of the mapping is split and all mapping is across three (3) zone types, as seen in Table 1.

Table 1. Map of ports according to *zone type* – scale-out from 4 to 6 grid controllers

Zone type name	Grid controller 1	Grid controller 2	Grid controller 3	Grid controller 4	Grid controller 5	Grid controller 6
	Ports 1 and 3	Ports 1 and 3	Ports 1 and 3	Ports 1 and 3	Ports 1 and 3	Ports 1 and 3
Z1234	✓					
Z3456			✓ Added			
Z1256	✓ Added				✓ Added	

iSCSI network configurations

Host systems can connect to a FlashSystem A9000R over an Ethernet network using the Internet Small Computer System Interface (iSCSI) protocol. Use these important recommendations when configuring your iSCSI host port connections, in order to achieve high availability and high performance in your storage system.

- **iSCSI host post configuration**
  Use this information to properly configure your iSCSI (Ethernet) host port connectivity.
- **iSCSI connectivity requirements**
  Use these connectivity requirements in order to ensure redundancy, further protecting your system data.
- **iSCSI best practices**
  Use this information for iSCSI (Ethernet) host connectivity best practices.
- **iSCSI connectivity during scale-out (MES)**
  Use this information for preparing host connectivity during system scale-out (MES).

iSCSI host post configuration

Use this information to properly configure your iSCSI (Ethernet) host port connectivity.

Disabling and enabling Spanning Tree Protocol (STP) or Rapid Spanning Tree Protocol (RSTP)

When setting up your Ethernet switch for host port connectivity, disable the STP or RSTP for each port that is connected to the FlashSystem A9000R system on the switch. Keeping these enabled can cause loss of host connectivity to the FlashSystem A9000R during hot upgrade.

Your network administrator can disable and enable the STP or RSTP at any time.

The following are examples of how to disable and enable an STP on a specific interface, when using a Cisco Nexus 5000 Series Switch.

To disable the STP

```
switch# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
switch(config)# interface Eth101/1/12
switch(config-if)# no spanning-tree port type edge trunk
switch(config-if)#
```

To enable the STP

```
switch# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
switch(config)# interface Eth101/1/12
switch(config-if)# spanning-tree port type edge trunk
Warning: Edge port type (portfast) should only be enabled on ports connected to a single
 host. Connecting hubs, concentrators, switches, bridges, etc... to this
 interface  when edge port type (portfast) is enabled, can cause temporary bridging loops.
 Use with CAUTION
switch(config-if)#
```

Physical host port configuration

Be sure to connect port 1 of each grid controller in the system to switch 1 and port 3 of each grid controller to switch 2.

Connectivity for high performance

For high performance, follow these important practices:

- Try to spread all ports of grid controllers in the system evenly.
- Utilize the CPU in each grid controller, as much as possible, for high CPU utilization. To achieve this, it is recommended to use ports 1 and 3 for host connectivity and ports 2 and 4 for migration and mirroring activities (see Migration and mirroring connectivity).

Ensuring high availability of each host

In order to ensure high availability of each host, follow these important practices:

- Divide the paths between the two subnets: half the paths connecting to subnet 1 and the other half connecting to subnet 2.
- If possible, spread the paths across different grid controllers.
- Use different adapters within each of the grid controllers, as shown in Figure 1.

Figure 1. Crossing Fibre Channel grid controller ports

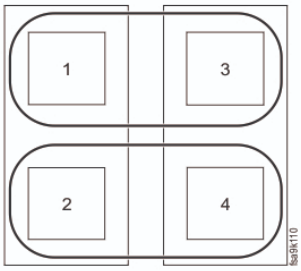

Important: It is important to use both Ethernet adapters for resiliency and high performance.
Figure 2 shows the port numbers for each of the iSCSI ports on the grid controllers.

Figure 2. iSCSI port numbering on the grid controllers

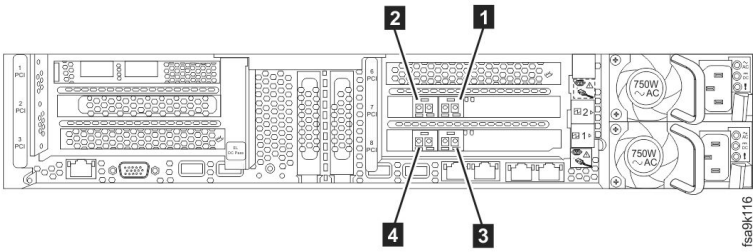

- 1 10 GbE port 1
- 2 10 GbE port 2
- 3 10 GbE port 3
- 4 10 GbE port 4

Only use non-redundant configurations when the risks of a single point of failure are acceptable, which is typically the case for test and development environments. Non-redundant configurations should generally not be used. For connectivity requirements, see iSCSI connectivity requirements.

iSCSI connectivity requirements

Use these connectivity requirements in order to ensure redundancy, further protecting your system data.

In a production environment, always connect Ethernet hosts to a minimum of two separate Ethernet switches to provide redundancy. In Ethernet network configurations, there are two dual-port connections per grid controller.
Note:

- For best performance, use a 10 Gb Ethernet switch and HBA on the host to obtain maximum performance or throughput.
- Host system can have as many HBAs as needed to support the operating system, application, and overall performance requirements.
- In a FlashSystem A9000R system with iSCSI configuration, each grid controller has two dual 10 Gb Ethernet ports.

Several network configurations using Ethernet are technically possible, and each configuration varies in terms of cost, flexibility, performance, and reliability.

In the FlashSystem A9000R system, each iSCSI port is defined with its own IP address. Before you set up the iSCSI network, gather the following information for each iSCSI port:

- IP address
- Net mask

- Default gateway
- Maximum transmission unit (MTU)

Note the following when setting up your iSCSI network:

- If two or more iSCSI connections exist in your iSCSI network, iSCSI ports for mirroring must be defined on switches with flow control set to on.
- MTU configuration is required if the network supports an MTU that is greater than the default 1536 bytes.
  The maximum MTU value supported is 9216 bytes.

- Because the FlashSystem A9000R system acts as a TCP server for iSCSI connections, packets are always routed through the Ethernet port from which the iSCSI connection was initiated. The default gateways are required only if the hosts are not on the same layer-2 subnet as the FlashSystem A9000R system.
- If present, Ethernet VLANs and IP routers must be configured to enable connectivity between the host systems and the FlashSystem A9000R system.
- If present, IP routers must be configured to enable access between the hosts and the FlashSystem A9000R system.

Important: Link aggregation is not supported. Ports cannot be bonded.
For supported interoperability configurations, see the System Storage® Interoperation Center website .

iSCSI best practices

Use this information for iSCSI (Ethernet) host connectivity best practices.

Talk to your IBM® representative to help optimally setup your host connections, according to your needs.

In order to achieve high performance, it is important to:

1. Spread the host connections to each grid controller evenly.
2. Utilize the CPU in each grid controller as much as possible; therefore, it is recommended to use ports 1 and 3 and then ports 2 and 4.

It is recommended to use half the ports in each grid controller (ports 1 and 3) and create three identical subnets in each switch.

Table 1 is an example of how the subnets should be divided in a full rack configuration.

Table 1. Example of recommended switch subnets

Model	Switches	Subnet	Grid controllers	Port numbers
415, 425, and U25	1, 2	1	Grid controllers 1, 2, 3, 4	1, 3
415, 425, and U25	1, 2	2	Grid controllers 5, 6, 7, 8	1, 3
415 only	1, 2	3	Grid controllers 9, 10, 11, 12	1, 3

This configuration creates high availability, high performance, and balancing. In addition, all grid controllers are utilized for each host.
See Table 2 for an example of how the host connections should be divided in model 425 and U25 minimum racks with three grid controllers.

Table 2. Example of recommended Ethernet host connections for model 425 and U25 minimum rack installations

Switch	Subnet	Grid controllers	Port number
1	1	Grid controllers 1, 2, 3	1
2	1	Grid controllers 1, 2, 3	3

iSCSI connectivity during scale-out (MES)

Use this information for preparing host connectivity during system scale-out (MES).

For more information about system scale-out, see Conducting system scale-out (MES).

- The original physical port connections are kept.
- Continue the physical port connection divisions, as described in iSCSI host post configuration.

Important: When reconfiguring subnets for scale-out situations, be sure that the hosts remain spread across the subnets evenly.

Network cable requirements

The customer is responsible for supplying cables that connect to the patch panel, including host attachment (Fibre Channel or iSCSI) cables, management cables, maintenance cables, and virtual private network (VPN) cables.

Table 1 shows the various cable type requirements for your system.

Table 1. Required cable types

Cables	Required cable type
Fibre Channel	50 µm (micrometer) multimode Fibre Channel cables with LC/LC or LC/SC connectors (if applicable). Note: If you require 62.5-µm fibers, contact an IBM® service representative for assistance.
Ethernet (iSCSI)	50 µm (micrometer) multimode optical cables with LC/LC or LC/SC connectors.
Management	Straight copper gigabit CAT5e-rated Ethernet cables with RJ-45 connectors.
Virtual private network (VPN)	Straight copper CAT5e-rated Ethernet cables with RJ-45 connectors.

Network and host connectivity security information

The storage system integrates various security features, to protect your network and host systems.

Use this information to help understand and plan for network and host connectivity security for your IBM FlashSystem® A9000R storage system.

- **Internet Protocol Security (IPSec)**
  Internet Protocol Security (IPSec) is a protocol suite that allows for enhanced security of IP communications through the authentication and encryption of IP packets.
- **Data-at-rest encryption**
  The storage system secures all written data with industry-standard AES encryption for data-at-rest.
- **User authentication and access control**
  Use this information to understand how to securely maintain your passwords.
- **PCI DSS compliance**
  The Payment Card Industry Data Security Standard (PCI DSS) is the global information security standard, for organizations that process, store, or transmit data with any of the major credit card brands. IBM FlashSystem A9000R systems comply with PCI DSS standards.

Internet Protocol Security (IPSec)

Internet Protocol Security (IPSec) is a protocol suite that allows for enhanced security of IP communications through the authentication and encryption of IP packets.

The IBM FlashSystem® A9000R system software and IBM® management tools allow for the use of passkey or certificate authentication to establish IPSec connectivity between management workstations and the management or VPN ports of the storage system.

Configuration of the IPSec is done through the command-line interface (CLI). For IPSec configuration, see the *IBM FlashSystem A9000R Command-Line Interface (CLI) Reference Guide* , SC27-8711 on the IBM FlashSystem A9000R documentation .

Data-at-rest encryption

The storage system secures all written data with industry-standard AES encryption for data-at-rest.

Encryption key management can be carried out through an external or an internal scheme.

Encryption can be enabled during the installation of the system or at any time later. While encryption is not enabled, the system might not meet customers or legal compliance standards and the data might not be protected against security issues. Encryption can be disabled only when no volumes are defined.

Encryption is configured through the CLI.

- For more information on data-at-rest encryption, see the *IBM FlashSystem® A9000R Product Overview* , SC27-8559.
- For information regarding encryption CLI commands, see the *IBM FlashSystem A9000R Command-Line Interface (CLI) Reference Guide* (SC27-8711) .

- **Internal data-at-rest encryption scheme**
  The internal encryption key management scheme generates and stores the encryption key locally, within the storage system.
- **External data-at-rest encryption scheme**
  An external encryption key management scheme stores the keys separately from the data, thereby presenting a secured and well-defined interface for key services.

Internal data-at-rest encryption scheme

The internal encryption key management scheme generates and stores the encryption key locally, within the storage system.

The default encryption for your FlashSystem A9000R system is the internal encryption key management scheme.

With the internal encryption key management scheme, keys are not affected by software upgrades and remain available upon the failure of up to two grid controllers.

Because the encryption is internal, it contains more security risks if the system is physically accessed. (For steps of preventing physical access to your system, see Site security considerations).

In addition, further precautions are necessary during relocation of your system. The boot drives contain the encryption keys and must be shipped separately for the security of your system. For specific instructions on relocating with an internal encryption key management scheme, see Relocation shipping requirements for systems with local encrypted schemes.

External data-at-rest encryption scheme

An external encryption key management scheme stores the keys separately from the data, thereby presenting a secured and well-defined interface for key services.

The separation of key storage from data storage and key management is accomplished with external Key Management Interoperability Protocol (KMIP) compliant servers, such as IBM® Security Key Lifecycle Manager (SKLM) or Gemalto SafeNet KeySecure server.

The separation of the keys from the data provide another layer of security should your system be physically accessed (see Site security considerations).
Note: To protect against the possibility that all Security Key Lifecycle Managers (SKLMs) become unusable and unrecoverable (for example, following a disaster, or other difficulties during the relocation process), the system enables you to create a *recovery key*. With a recovery key, Security Administrators can unlock an IBM FlashSystem A9000R system without the involvement of a key server.
For more information regarding recovery keys and how to use them, see the Redbook publication Data-at-rest Encryption for the IBM Spectrum Accelerate Family: IBM XIV and IBM FlashSystem A9000 and A9000R .

For more information, and purchasing options, speak to your IBM representative.

User authentication and access control

Use this information to understand how to securely maintain your passwords.

The storage system features role-based authentication, either natively or through use of LDAP-based authentication.
Attention:

- Be careful when saving user credentials locally. This information can be accessed if site or system access is obtained.
  For more information, see Site security considerations.
- To properly ensure password safety, be sure to change passwords often to limit security risks.
- For more information on mirroring, see:
  - Migration and mirroring connectivity
  - *IBM FlashSystem A9000R Product Overview* , SC27-8559 on the IBM FlashSystem A9000R documentation
- For more information on user authentication and access control, see the *IBM FlashSystem A9000R Product Overview* , SC27-8559 on the IBM FlashSystem A9000R documentation
- **Native user authentication**
  Native user authentication is the default mode for authenticating users and user groups.
- **External authentication via Lightweight Directory Access Protocol (LDAP)**
  IBM FlashSystem A9000R systems offer the capability to use LDAP server-based user authentication.

Native user authentication

Native user authentication is the default mode for authenticating users and user groups.

In this mode, users and groups are authenticated against a database locally on the system. The authentication is based on the submitted username and password, which are compared to user credentials defined and stored on the storage system.

The authenticated user must be associated with a user role that specifies the system access rights.
Note: Technicians use a different form of native user authentication. For more information regarding technician authentication, see Support and software maintenance security information.

External authentication via Lightweight Directory Access Protocol (LDAP)

IBM FlashSystem® A9000R systems offer the capability to use LDAP server-based user authentication.

Note:

- For supported LDAP server products, see Product selection.
- For more detailed information about LDAP products, role mapping, defining on your storage system, and more, speak to your IBM representative and refer to:
  - *IBM Hyper-Scale Manager for IBM Spectrum Accelerate Family* Redbook, SG24-8376
  - *Enabling LDAP for IBM FlashSystem A9000 and A9000R* with Microsoft Active Directory Redpaper

There are three user-authentication options:

LDAP only
: Requires that all users must be granted LDAP credentials in order to gain access. This mode does not allow concurrent native user authentication.

Native user authentication (internal users)
: Grants access only to internal users (such as technicians, Admin users, and so on) who do not use LDAP authentication.

Mixed
: Allows concurrent use of LDAP authentication and native-user authentication protocols.

Important: As a preferred practice, the LDAP server and the FlashSystem A9000R storage system should have their clocks synchronized to the same time source, be registered, and be configured to use the same DNS servers.
LDAP authentication must be set up separately, as the native user authentication mode (for internal users) is the default. For more information about native user authentication, see Native user authentication.

Managing multiple systems in LDAP authentication mode and single sign-on (SSO)

The task of managing multiple IBM FlashSystem A9000 and A9000R systems can be simplified by using LDAP authentication mode.

As a result of all user credentials being stored centrally in the LDAP directory, it is no longer necessary to synchronize user credentials among multiple storage systems. After a user account is registered in LDAP, multiple storage systems can use credentials stored in the LDAP directory for authentication.
Note: LDAPs are all located externally, posing less risk should the system be physically accessed by unwanted persons.
For more information about the physical security of your system, see Site security considerations.

Because the user's password is stored in the LDAP directory, all connected storage systems authenticate the user with the same password. If the password is changed, all storage systems automatically accept the new password.

This mode of operation is often referred to as *single sign-on* (SSO). This approach is especially useful in remote mirroring configurations, where the storage administrator is required to frequently switch from source to target system. For more information about remote mirroring, see the IBM FlashSystem A9000R Product Overview, SC27-8559 on the IBM FlashSystem A9000R documentation .

Product selection

LDAP authentication of the storage system supports three LDAP server products:

- Microsoft Active Directory
- Oracle Directory Server Enterprise Edition
- OpenLDAP

The current skill set of your IT staff is always an important consideration when choosing any products for centralized user authentication. If you have skills in running a particular directory server, it might be a wise choice to standardize on this server because your skilled people will best be able to customize and tune the server. Your experts will be able to provide the most reliable and highly available implementation for the LDAP infrastructure.

Security LDAP with Secure Sockets Layer (SSL)

In any authentication scenario, information is exchanged between the LDAP server and your storage system where access is being sought. SSL can be used to implement secure communications between the LDAP client and server. LDAP over SSL (LDAPS), the secure version of the LDAP protocol, allows a setup where user passwords are never transferred in clear text.

SSL provides methods for establishing identity using X.509 certificates and ensuring message privacy and integrity using encryption.

To create an SSL connection, the LDAP server must have a digital certificate signed by a trusted certificate authority (CA). Companies have the choice of using a trusted CA from another vendor or creating their own certificate authority.

To be operational, SSL must be configured on both the client (IBM FlashSystem A9000 or FlashSystem A9000R storage system) and the LDAP server. Server configuration includes generating a certificate request, obtaining a server certificate from a CA, and installing the server and CA certificates.
Important: When defining the LDAP server with a security certificate in the storage system, the fully qualified name of the LDAP server must match the "issued to" name in the client's certificate.
LDAP registration in the server can be done either through the UI or through the CLI. LDAP registration through the UI can be done with a simple file upload. When using the CLI, Windows users can drag and drop the certificate into the XCLI utility window. Other uses need to copy and paste a long string containing the certificate, noting the correct formatting.
Maintaining SSL certificates

New SSL certificates must be installed before the existing ones expire.

PCI DSS compliance

The Payment Card Industry Data Security Standard (PCI DSS) is the global information security standard, for organizations that process, store, or transmit data with any of the major credit card brands. IBM FlashSystem® A9000R systems comply with PCI DSS standards.

Table 1 describes how IBM FlashSystem A9000R complies with these standards.

Table 1. PCI-DSS Support

Requirement	PCI-DSS Section	FlashSystem A9000R solution
Encrypt all non-console administrative access	2.3	All management connections are secured via IPSec.
Implement a data retention and disposal policy that includes...Processes for secure deletion of data when no longer needed	3.1.1	FlashSystem A9000R provides data-at-rest encryption by use of SED capabilities of the flash enclosure IBM® MicroLatency® modules and by encrypting the SSD vault devices.
Disk encryption and key management requirements	3.4.1, 3.5, 3.6	Key management using IBM's SKLM key server services, using KMIP key exchange protocol. Disks are encrypted using AES256 in XTS mode.
Render all passwords unreadable during transmission and storage on all system components using strong cryptography	8.4	XIV stores secure and irreversible hashes of user passwords

Requirement	PCI-DSS Section	FlashSystem A9000R solution
Enable accounts used by vendors for remote access only during the time period needed. Monitor vendor remote access accounts when in use.	8.5.6	**xiv_support_enable/disable** is used to temporarily enable remote access
Change user passwords at least every 90 days	8.5.9	Enforcement of password expiration may be provided using LDAP servers, as configured by the system administrator. Note: Password expiration rules do not apply to the storage admin user.
Minimum password length ... passwords containing both numeric and alphabetic characters ... Limit repeated access attempts ... Set the lockout duration to a minimum of 30 minutes	8.5.10 - .14	Enforcement of password rules may be provided using LDAP servers, as configured by the system administrator. Note: Password enforcement rules do not apply to the storage admin user.
If a session has been idle for more than 15 minutes, require the user to re-authenticate	8.5.15	Supported by IBM Hyper-Scale Manager UI and XCLI utility.
Audit trails	10.5.1- .7	The audit trails are supported through the syslog (Service Center) server.

Migration and mirroring connectivity

Use this information in order to plan for physical connectivity for data migration and mirroring.

This section covers only the important physical connection setup required for migration and mirroring. For more in-depth information, see the following relevant publications:

- *IBM FlashSystem® A9000R Product Overview* , SC27-8559
- IBM® FlashSystem A9000 and A9000R Business Continuity Functions Redbook, REDP-5401

- **Physical connectivity for mirroring and migration using Fibre Channel**
  Be sure to use ports 2 and 4, 4 being the initiator, when conducting mirroring or migration using FC connectivity.
- **Physical connectivity for mirroring and migration using iSCSI**
  Port connectivity for mirroring and migration using the Ethernet ports depends on the grid controller type being used in the system.
- **Migration and mirroring best practices**
  Use this information in order to achieve high availability and high performance in your storage system when performing migration and mirroring.

Physical connectivity for mirroring and migration using Fibre Channel

Be sure to use ports 2 and 4, 4 being the initiator, when conducting mirroring or migration using FC connectivity.

Ports 1 and 3 are used for host connectivity. For more information, see Host network connectivity and multipathing configurations.
For more information about FC adapters, see Fibre Channel (FC) adapter types.

shows the port numbers for each of the FC ports on the grid controllers.

Figure 1. FC port numbering on the grid controllers

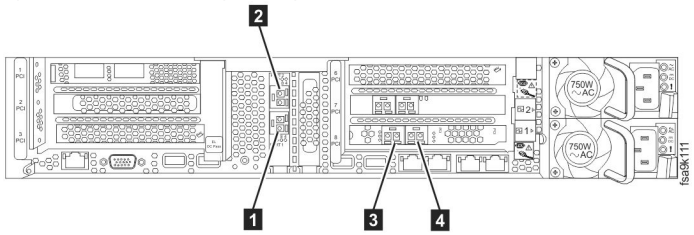

- 1 FC port 1
- 2 FC port 2
- 3 FC port 3
- 4 FC port 4

For more information about grid controller feature codes and connectivity options, see Grid controller components and feature codes.

Physical connectivity for mirroring and migration using iSCSI

Port connectivity for mirroring and migration using the Ethernet ports depends on the grid controller type being used in the system.

Ports 1 and 3 are used for host connectivity. For more information, see Host network connectivity and multipathing configurations.

iSCSI on grid controllers containing FC ports (feature codes 5001, 5003, 5005)

Important: In this grid controller setup, migration can **only** be done via FC port connections (see Physical connectivity for mirroring and migration using Fibre Channel).
When conducting mirroring using Ethernet port connectivity, use ports 2 and 4 (where applicable).
Ports 1 and 3 are used for host connectivity. For more information, see Host network connectivity and multipathing configurations.

Figure 1 shows the port numbers for each of the Ethernet ports on the grid controllers.

Figure 1. Ethernet (10 Gb) port numbering on the grid controllers (mixed FC and iSCSI configuration)

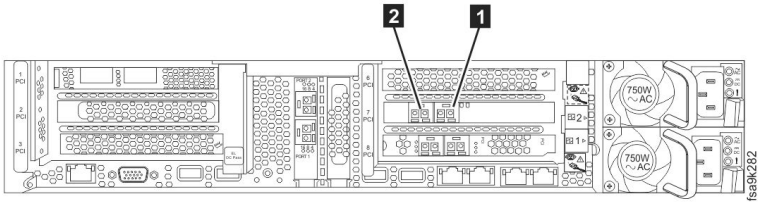

- 1 10GbE port 1
- 2 10GbE port 2

iSCSI on grid controllers containing full iSCSI connectivity (feature codes 5002, 5004)

When conducting mirroring or migration using Ethernet port connectivity, use ports 2 and 4.

Figure 2 shows the port numbers for each of the Ethernet ports on the grid controllers.

Figure 2. Ethernet (10 Gb) port numbering on the grid controllers

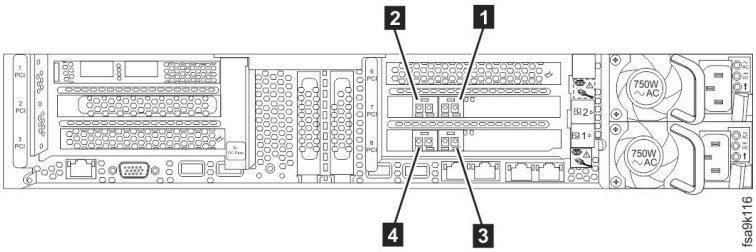

- 1 10GbE port 1
- 2 10GbE port 2
- 3 10GbE port 3
- 4 10GbE port 4

Migration and mirroring best practices

Use this information in order to achieve high availability and high performance in your storage system when performing migration and mirroring.

Talk to your IBM® representative to help optimally setup any migration and mirroring connectivity, according to your needs.

To achieve high availability and performance, it is important to:

- Spread all ports being used for migration or mirroring connectivity over multiple grid controllers, evenly distributing the mirroring load.
- Use both Fibre Channel or iSCSI adapters within each of the grid controllers, as shown in Figure 1.

Figure 1. Crossing grid controller adapter ports

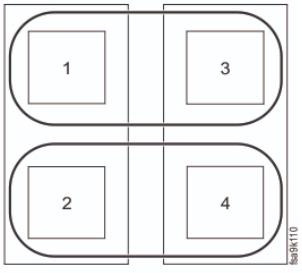

- Divide the paths between the two switches or subnets: half the paths connecting to switch/subnet 1 and the other half connecting to switch/subnet 2. When possible, P2 connects to switch/subnet 1 and P4 connects to switch/subnet 2.

The administrator must ensure that:

- Host connections avoid single points of failure by applying redundant connections.
- All host workload is adequately balanced across the connections and grid controllers, ensuring system resource utilization is maximized.

Be sure to follow these **minimum** connectivity requirements:

Note: For more information on grid elements, see FlashSystem A9000R grid elements

Fibre Channel connections in systems with two or more grid elements
: Use a minimum of three FC ports configured as *initiator* and an additional three FC ports configured as *target* on each system. Each pair of FC ports (one initiator and one target) must reside on separate grid controllers, in order to provide redundancy and sustain two grid controller failures.
Using additional FC ports helps improve performance.

Fibre Channel connections in model 425 and U25 minimum configurations (3 grid controllers)
: Use a minimum of two FC ports configured as *initiator* and an additional two FC ports configured as *target* on each system. Each pair of FC ports (one initiator and one target) must reside on separate grid controllers, in order to provide redundancy and sustain a grid controller failure.
Using additional FC ports helps improve performance.

Ethernet connections in systems with two or more grid elements
: Use a minimum of three bi-directional Ethernet connections, using three iSCSI ports on each system. Each Ethernet connection must reside on separate grid controllers, in order to provide redundancy and sustain two grid controller failures.
Using additional bi-connections on iSCSI ports helps improve performance.

Ethernet connections in model 425 and U25 minimum configurations (3 grid controllers)
: Use a minimum of two bi-directional Ethernet connections, using two iSCSI ports on each system. Each Ethernet connection must reside on a separate grid controller, in order to provide redundancy and sustain a grid controller failure.
Using additional bi-connections iSCSI ports helps improve performance.

Planning for physical shipment

Ensure that your environment meets the standard delivery clearance and weight requirements for the IBM FlashSystem® A9000R system.

Prepare for equipment delivery so that professional movers or riggers can transport the equipment to the final installation site. If you cannot complete preparations at the time of delivery, you must make your own arrangements for the professional movers to complete transportation later.
Important: Use only professional movers to transport the equipment.
The IBM® service representative can minimally reposition the rack at the installation site, as required.

- **Planning to receive delivery**
  The professional movers or riggers are responsible for delivering and unloading the IBM FlashSystem A9000R system as close to its final destination as possible. You must ensure that the loading ramp and receiving area can accommodate the storage system shipment.
- **Planning for relocation**
  When unpacking your storage system delivery, be sure to keep the original packaging material, in case it is needed for relocation at a later date.
- **Shipment weights and dimensions**
  To help you plan for the delivery of your FlashSystem A9000R, ensure that the loading dock and receiving area can support the weight and dimensions of the packaged FlashSystem A9000R shipments.

Planning to receive delivery

The professional movers or riggers are responsible for delivering and unloading the IBM FlashSystem® A9000R system as close to its final destination as possible. You must ensure that the loading ramp and receiving area can accommodate the storage system shipment.

Procedure

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

1. Coordinate a technical survey with the IBM® service representative to plan the loading path from the truck to the server room.
2. Determine the packaged weight and dimensions of the FlashSystem A9000R system container and other containers that you to be received (see Shipment weights and dimensions).
3. Ensure that the loading dock, receiving area, all doors, and elevators can safely support the packaged weight and dimensions of the shipping containers.
   If the rack height or weight must be reduced for delivery to locations where doorway heights are smaller than the delivery clearances, the height and weight reduction shipping options (feature codes AFR2 or AFR3) must be ordered. For information about the height and weight reduction feature codes, see Height reduced shipping option and Weight reduced shipping option.
   Important: The rack is more easily moved when on a pallet, and the rack casters might damage floors and carpets; therefore, roll the rack into position on its own casters only when necessary.
4. Ensure that there is a clear and level path from the truck to the building entrance. Ensure that there are no steps from the truck to the installation site.
5. Ensure that the loading ramp at your site does not exceed an angle of 10 degrees, as shown in the following figure.

   Figure 1. Maximum tilt for a packaged rack is 10 degrees

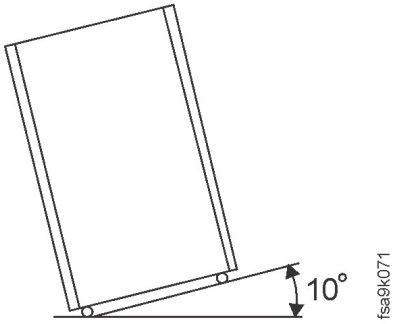

A ramp with a maximum angle of 10 degrees must not be higher than 5.4 cm (2.1 in.) vertically for each 30.5 cm (12 in.) of horizontal length.

Planning for relocation

When unpacking your storage system delivery, be sure to keep the original packaging material, in case it is needed for relocation at a later date.

Important: Whenever IBM FlashSystem® A9000R units need to be physically moved to another location, the relocation must only be performed by an IBM service provider.

Customer responsibilities
When relocating systems, the customer is responsible for the following:

- Informing their IBM representative what packing requirements are necessary for their relocation (based off of machine type, model, and encryption type)
- **For locally encrypted systems only:** Shipping the drive packages

IBM responsibilities
When relocating systems, IBM is responsible for the following:

- Ordering the packaging material for the customer (at customer expense)
- **For locally encrypted systems only:** Packaging and labeling the location of the boot drives
- **For locally encrypted systems only:** Reinstalling the boot drives at the new customer location

- **Relocation shipping requirements for systems with external encryption schemes**
  IBM FlashSystem A9000R systems that are encrypted with external key management schemes must have access to key servers with the original key upon restart.
- **Relocation shipping requirements for systems with local encrypted schemes**
  Shipping systems that contain encrypted data, with the encryption keys, outside of secure customer premises, increases the risk of confidential data exposure.

Relocation shipping requirements for systems with external encryption schemes

IBM FlashSystem® A9000R systems that are encrypted with external key management schemes must have access to key servers with the original key upon restart.

About this task

Only use this procedure for systems that are encrypted with external key management schemes. For systems encrypted with local key management schemes, see Relocation shipping requirements for systems with local encrypted schemes.

Procedure

Use the following procedures on systems using external key encryption before relocation.

1. Ensure that you have any new IP addresses needed for the new location (ie. system IP, network component IPs, SMTP GW, and any other IP addresses).
2. Give the new IP address to your service technician before system shutdown, for use when connecting the system upon relocation.
   Important: If these steps are not taken for externally encrypted systems, the system will not be able to restart.
   To protect against the possibility that all Security Key Lifecycle Managers (SKLMs) become unusable and unrecoverable (for example, following a disaster, or other difficulties during the relocation process), the system enables you to create a *recovery key*. With a recovery key, Security Administrators can unlock an FlashSystem A9000R system without the involvement of a key server.

   For more information regarding recovery keys and how to use them, see the Redbooks publication Data-at-rest Encryption for the IBM Spectrum Accelerate Family: IBM XIV and FlashSystem A9000 and A9000R.

Relocation shipping requirements for systems with local encrypted schemes

Shipping systems that contain encrypted data, with the encryption keys, outside of secure customer premises, increases the risk of confidential data exposure.

To ensure the safety of your system, it is strongly recommended to perform the following during relocation of locally encrypted systems:

- Separate the data (system) and the encryption keys (boot media devices) prior to relocation. This action can only be performed by IBM.
- Ship the boot media devices separately from the rest of the data.
- Ship the boot media devices in separate packages.
  The shipments should be a *minimum* of two (2) separate packages.
  Note: In order to avoid risk of loss of access to system data, all drives labeled 0 should be shipped in a separate shipment and using a different carrier than drives labeled 1.

Shipment weights and dimensions

To help you plan for the delivery of your FlashSystem A9000R, ensure that the loading dock and receiving area can support the weight and dimensions of the packaged FlashSystem A9000R shipments.

At least one shipping container is delivered for each FlashSystem A9000R that you order. The container is a wooden pallet that is covered by a corrugated fiberboard (cardboard). This container contains the FlashSystem A9000R system rack with components installed, and other items such as, power cords, CDs and printed publications, and other features or peripherals for your model. Additional containers are delivered for optional features such as the rear-door heat exchanger.

Delivery clearance requirements

The clearance measurements (height × width × depth) that are required for delivery through all doors and elevators are listed in the following table. These measurements are for the typical height of the pallet. The measurements do not include more clearance that is needed to raise the pallet on a pallet jack for movement.

Table 1. Typical delivery clearance requirements

Dimension	Clearance requirement
Height	216 cm (85.0 in.)
Depth	144 cm (56.7 in.)
Width	94 cm (37.0 in.)

Delivery weight requirements

The path from the truck and to the server room must support the weight of the rack, including packaging materials.

Table 2. Floor weight-support requirements

Grid element configuration	Total weight
12 grid controllers and 6 flash enclosures (model 415 only)	969 kg (2136 lbs)
10 grid controllers and 5 flash enclosures (model 415 only)	889 kg (1960 lbs)
8 grid controllers and 4 flash enclosures	809 kg (1783 lbs)
6 grid controllers and 3 flash enclosures	729 kg (1607 lbs)
4 grid controllers and 2 flash enclosures	649 kg (1430 lbs)
3 grid controllers and 1 flash enclosure (models 425 and U25 only)	578 kg (1274)

Planning for remote support, on-site service, and maintenance

This section provides you with information to help prepare you for optimal support and software maintenance.

- **Planning for remote support connection**
  IBM® Remote Support Center (RSC) is a management system used by authorized IBM service representatives to provide remote support and problem-determination assistance over a secure network interface.
- **Planning for Call Home**
  Using Call Home, you can set up the IBM FlashSystem® A9000R system to automatically send pre-failure or failure notifications to the IBM Troubleshooting Ticketing System in the IBM Service Center. You can also configure the storage system to automatically send alerts directly to you.
- **Required support information**
  Use this information to help prepare for support requests.
- **Support and software maintenance security information**
  The IBM Remote Support Center (RSC) provides a high level of security for a remotely supported FlashSystem A9000R system through encryption, authentication, authorization, auditing, and field-proven security components.

Planning for remote support connection

IBM® Remote Support Center (RSC) is a management system used by authorized IBM service representatives to provide remote support and problem-determination assistance over a secure network interface.

The remote support connection is used either before or after parts are called out by the internal diagnostic tests of the storage system. IBM authorized personnel use remote support connections to perform real-time problem analysis and isolation.

Remote support access is obtained over a dedicated network connection. When using this method, the storage system is connected to the RSC through one of the virtual private network (VPN) ports or management ports on the patch panel. For more connectivity information, see Network connections for management.

The RSC has three components:

Software that is installed on the storage system and handles remote support connectivity.
: It relies on a single outgoing TCP connection and is not able to receive inbound connections of any kind. The remote support client is controlled by using the command-line interface (CLI) commands and starts a connection, terminates a connection (due to timeout or customer request), and attempts to reconnect when the connection is terminated unexpectedly.

Front servers that serve as a hub at which the storage system and the remote-support back server connect.
: The front servers are located in an IBM demilitarized zone (DMZ) and receive and maintain connections from the remote support client and the back server. The front servers are strictly inbound and do not initiate any outbound communication.
No sensitive information is stored on the front server, and all data passing through the front server from the client to the back server is encrypted, so the front server or a malicious entity in control of a front server cannot access this data.

One or more *back servers* are located within the IBM intranet.
: Only IBM service representatives that are authorized to perform remote support of the storage system can access these servers.
The back server authenticates the IBM service representative, provides the IBM service representative with a user interface through which to choose a system to support, and manages the remote support session as it progresses. The IBM service representative connects to the back server by using a Secure Shell (SSH) client or an HTTPS connection with any browser.

Requirement: To perform remote support through the Remote Support Center, your storage system must be able to initiate an outbound SSH connection to IBM. If the system does not have direct access to the Internet (for example, due to a firewall), you can use the IBM Remote Support Proxy (RSP) to facilitate the connection to IBM. For more information, see the *IBM XIV® Remote Support Proxy User Guide* (GA32-0795).
Figure 1 illustrates the remote support components.

Figure 1. Remote support components

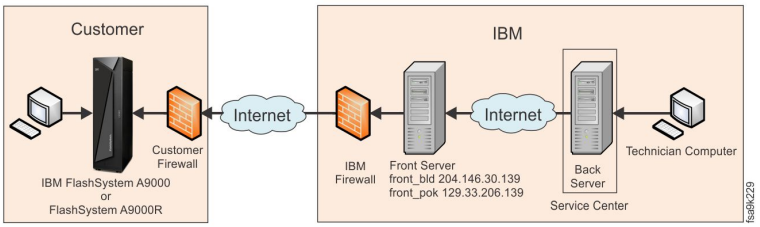

The RSC uses the Secure Shell (SSH) protocol for transporting data. The encryption used by SSH provides confidentiality and integrity of the transferred data even over insecure mediums.

To conduct a remote support session, the IBM service representative must explicitly connect to the RSC back server. The following secure remote connection process occurs when a support session is opened:

1. The customer initiates an Internet SSH connection to the RSC.
2. The RSC identifies the storage system and marks it as connected.
3. The IBM service representative connects to the RSC using SSH.
4. The RSC authenticates the IBM service representative.
5. The IBM service representative is shown a list of currently connected storage systems that correspond with the defined permissions, or the IBM service representative manually enters the serial number to view the system.
6. The IBM service representative chooses the storage system to support. Only permitted systems are displayed, and all activity is logged.
7. The fully recorded support session commences.
8. The IBM service representative terminates the support session.
9. The system disconnects from the RSC.

While a support session is in progress, the storage system displays the system status on the IBM Hyper-Scale Manager user interface (UI). You can view the process of phasing out a component or the restart of customer-visible system services as it happens. The customer has full control over whether to proceed with a support session by using mechanisms such as timeout or force-disconnect. If a session disconnects unexpectedly, the IBM service representative can resume the session when the storage system next connects to the RSC.

- **Remote support for severe system conditions**
  A remote support connection may be established automatically, when activated, if severe system conditions (critical issues) are seen in the system and access to the machine is blocked to the host. Immediate action is necessary so that waiting for session connection is not necessary. This function is also known as "XRSC on severe system conditions."

Remote support for severe system conditions

A remote support connection may be established automatically, when activated, if severe system conditions (critical issues) are seen in the system and access to the machine is blocked to the host. Immediate action is necessary so that waiting for session connection is not necessary. This function is also known as "XRSC on severe system conditions."

Important: Enabling the remote support for severe system conditions is strongly recommended in order to keep system repair time to a minimum and in order to resume connectivity to hosts as fast as possible. This allows an IBM service representative to access the system remotely and start the repair action immediately.
Note: XRSC on severe system conditions (XRSC) is compliant with General Data Protection Regulation (GDPR) EU 2016/679. For more information, refer to http://ibm.biz/gdpr-storage-paper.
Remote support for severe system conditions permits a remote support access without the need for the customer to initiate the SSH session towards the RSC. It also eliminates the need to dispatch an IBM service representative to the customer site in order to initiate this session to do so.
The remote support for severe system conditions can be configured by the IBM service representative at the time of the storage system installation. This information is communicated through the Technical and Delivery Assessment (TDA) checklist and worksheets.

If automatic remote support activation is not enabled at the time of installation, you can enable these actions at any time, using the XCLI utility. Use the **support_center_config** command to configure the automatic connection to a support center, setting **automatically_connect** to **yes**.

For more detailed information, see IP configuration commands > Configuring the support center connection to enable automatic connect on restart in the *IBM FlashSystem® A9000R Command-Line Interface (CLI) Reference Guide* (SC27-8711).

Planning for Call Home

Using Call Home, you can set up the IBM FlashSystem A9000R system to automatically send pre-failure or failure notifications to the IBM® Troubleshooting Ticketing System in the IBM Service Center. You can also configure the storage system to automatically send alerts directly to you.

Note: Call Home is compliant with General Data Protection Regulation (GDPR) EU 2016/679. For more information, refer to http://ibm.biz/gdpr-storage-paper.
Call Home cannot accept incoming communication, which means that the IBM Service Center cannot contact the FlashSystem A9000R system using Call Home. When certain events occur in the FlashSystem A9000R system, Call Home sends a notification to the IBM Service Center. After receiving the notification, IBM service personnel analyze the problem promptly and take appropriate action. If the problem requires service, an IBM service representative is sent to your site with any necessary replacement parts. With access to the FlashSystem A9000R system, IBM service personnel can perform service tasks, such as viewing error logs and problem logs or initiating trace and dump retrievals.

Call Home proactively reduces problem handling efforts and provides more efficient self-service solutions. This ensures success when using IBM products and services, reducing critical situation events. This will also help you save time and money while maintaining your IT environment.

Call Home notifications are sent through email from an SMTP server. You must have an SMTP email system available that the FlashSystem A9000R system can use to send outgoing Call Home emails to IBM.

If required, the customer email gateway can be configured to send Call Home information to IBM only via a secured channel. For more information, see Monitoring and Troubleshooting > Encrypting Call Home and heartbeat notifications in IBM FlashSystem A9000 and IBM FlashSystem A9000R: Architecture, Implementation and Usage Redbook, SG24-8345 on the IBM Redbooks® website .

Call Home information is configured by the IBM service representative at the time of storage system installation. This information is communicated through the Technical and Delivery Assessment (TDA) checklist and worksheets.

Customer responsibility

As the customer, it is your responsibility to configure the SMTP email system to enable the FlashSystem A9000R system to send outgoing emails for the Call Home function. The email configuration rules must not inhibit call home emails from being sent in real time. Emails with "Importance" or "X-Priority" in the email header must not be delayed.

For example, FlashSystem A9000R emails must not be placed in a queue for later delivery or filtered for priority. The administrator must verify correct configuration and function of the email system before installation. Failure to verify the email system might delay the successful installation or repair of the FlashSystem A9000R system.

Table 1 lists requirements that must be met to configure the Call Home function.

Table 1. Call Home configuration information

Call Home configuration requirement	Value	Comment
Customer SMTP port	25	The customer SMTP server must be reachable on port 25 from the customer-provided management IP address.
Customer SMTP address	IP address	This IP address is supplied by the customer.
Email destination address	xiv-callhome-eastern-hemisphere@vnet.ibm.com or xiv-callhome-western-hemisphere@vnet.ibm.com	Customer SMTP server must allow relaying to the following IBM email addresses based on geographical location: • **East:** EMEA, Asia, Australia, Africa, and the rest of the world • **West:** USA, Canada, Latin America, and the Caribbean Islands
Email source address	xiv@il.ibm.com or customer-defined	The default email address from which the email is sent. This email address can be customer-defined to conform to the customer email relay rules.

- **Planning for Call Home Web**
  Call Home Web allows you to view online support information for one or more IBM storage systems in a consistent and consolidated manner.

Planning for Call Home Web

Call Home Web allows you to view online support information for one or more IBM® storage systems in a consistent and consolidated manner.

The standard Call Home option (see Planning for Call Home) allows you to set up the system to automatically send pre-failure or failure notifications to the IBM Troubleshooting Ticketing System in the IBM Service Center. This information can help IBM Support to better diagnose issues with your systems and proactively assist in identifying them and in developing an action plan for their fast resolution.

While the standard Call Home information is accessible to IBM Support only, you can now view the information online through Call Home Web on the IBM Support Portal . The information is presented in a consistent and consolidated manner. It includes descriptive events, generated by the standard Call Home, an exportable system summary, software levels, and your system information.

Additional features allow you to view such details as the last automated problem report received, maintenance contract expiration date, last inventory received, and last heartbeat received. When configured to send notifications, Call Home web notifies you via email upon receipt of an event from Call Home.

Figure 1 shows the flow between IBM Call Home and Call Home Web.

Figure 1. IBM Call Home and Call Home Web

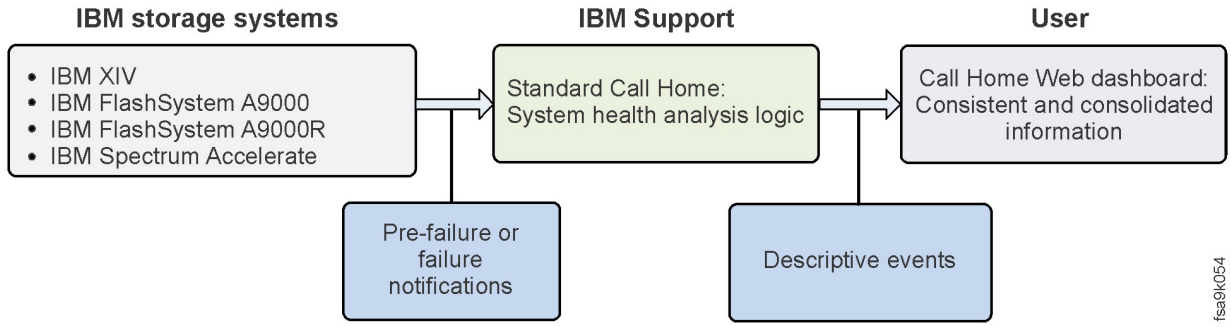

Attention: To use Call Home Web, you must first enable the standard Call Home option on your system, as explained in Planning for Call Home.
Note: Call Home Web is only available for systems under warranty or maintenance contract.
After Call Home has been enabled, you can add your system to Call Home Web. This operation is authorized for the following FlashSystem A9000R user roles:

- Storage integration administrator
- Storage administrator
- Application administrator

To add your system to Call Home Web, one of the above storage system users must issue the custom event similar to the following, but with properly substituted information values:

```
custom_event custom_event_description= "Connect System to Web Call Home" [severity = <INFORMATIONAL ]
Contact Name = 'Your Name': Contact Phone Number = '555-555-5555':
Contact Email Address = 'jdoe@company_name.com': ICN= '1234567'"
```

Upon receiving this event, the IBM Service Center generates a test event and establishes an association with your system.

To disconnect a storage system from Call Home Web:

- Go to the System Details page and click Remove System from Call Home Web.

Note: You can only remove a system that is fully added (activated and confirmed). If you need to remove a system that was accidentally added or cannot be confirmed, click Feedback and request that the system be removed.
For comprehensive information about Call Home Web, see the Call Home Web tab on the IBM Support Site assistance webpage .

Required support information

Use this information to help prepare for support requests.

Have the following information on hand when calling for a support request for a storage system:

- Phone number of machine location
- Machine serial number

Figure 1 depicts the location of the serial number on the front of a FlashSystem A9000R rack.

Figure 2 depicts the location of the serial number on the rear of a FlashSystem A9000R rack.

Figure 1. Machine type and model, and serial number label on front of rack

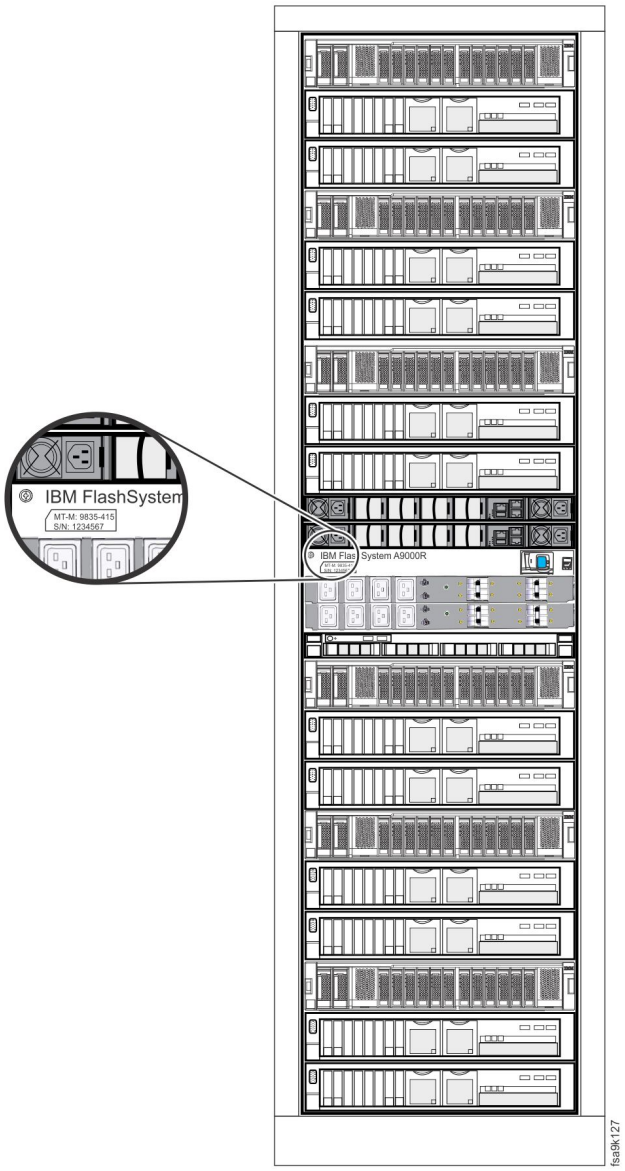

Figure 2. Machine type and model, and serial number on rear of rack

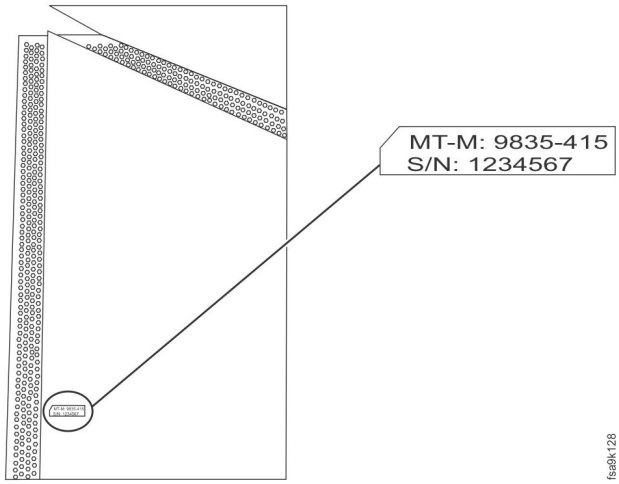

Support and software maintenance security information

The IBM® Remote Support Center (RSC) provides a high level of security for a remotely supported FlashSystem A9000R system through encryption, authentication, authorization, auditing, and field-proven security components.

On-site service security

Service representatives need a valid time certificate in order to connect to any given system.

Native user authentication

To prevent unauthorized access to the configuration of the storage system and ultimately to the information stored on its volumes, the IBM FlashSystem® A9000R storage system uses various forms of user authentication.

Customer user authentication
: Customer's use password-based user authentication.
Note: For more information on customer user authentication, see User authentication and access control.

Technician user authentication
: Service technicians have a challenge-response authentication protocol, establishing the authenticity of the technician.
This is done with a question (challenge) and comparing the answer (response) with information stored in a credential repository.

Local credential repository

By default, the IBM FlashSystem A9000R storage system is configured to use native (FlashSystem A9000R system managed) user authentication. Native user authentication uses the credential repository stored locally on the storage system. The FlashSystem A9000R local credential repository maintains the following information:

- User name
- User password
- User role
- User group
- Optional account attributes (such as email and phone numbers)

For more information about these role definitions and how to use them, see Security > Native user authentication > Local credential repository in IBM FlashSystem A9000 and IBM FlashSystem A9000R: Architecture, Implementation and Usage.

Deploying IBM® HyperSwap Quorum Witness

The following sections describe how to install, manage, and troubleshoot IBM Spectrum Accelerate Family HyperSwap® Quorum Witness.

Introduction

The IBM Spectrum Accelerate Family HyperSwap® Quorum Witness is a software component that facilitates coordination between two IBM FlashSystem A9000 or A9000R storage systems in a HyperSwap relation.

Note: For your convenience, IBM Spectrum Accelerate Family HyperSwap Quorum Witness is referred to as Quorum Witness throughout this user guide.
IBM HyperSwap provides high availability of host applications running on IBM FlashSystem A9000 and A9000R storage systems. It is designed to guarantee business continuity in case of hardware, power, or connectivity failure, or if a natural or man-made disaster occurs. To enable a non-disruptive automatic failover, the IBM HyperSwap solution includes Quorum Witness software, which must be installed at a separate site (failure domain 3).
Important: Quorum Witness cannot reside on the same set of FlashSystem A9000 and/or A9000R systems of which it serves as part of a HyperSwap configuration.
Note: Full description of the IBM® HyperSwap solution appears in the Product Overview of the IBM FlashSystem® A9000 and A9000R storage systems.
After the Quorum Witness software is installed on a dedicated server or a virtual machine, it is used by the connected IBM FlashSystem A9000 and A9000R storage systems as part of the HyperSwap solution. The Quorum Witness role in the solution is to track and report the systems health, using the storage systems' heartbeat messages, and serve as a tie breaker in split-brain scenarios. If a failure is detected, it resolves a potential contention between the peer storage systems, determining which of them will own the primary volume. The Quorum Witness is critical for complete high-availability functionality. Without a functioning Quorum Witness, automatic failover cannot be performed.

In the solution example illustrated below, data is replicated between two storage systems. The host is connected to a HyperSwap volume through both systems in an active-active configuration. The Quorum Witness is installed at a separate site (failure domain 3) and monitors the storage system health, by receiving heartbeat messages from both storage systems.

Figure 1. Typical HyperSwap solution with the Quorum Witness

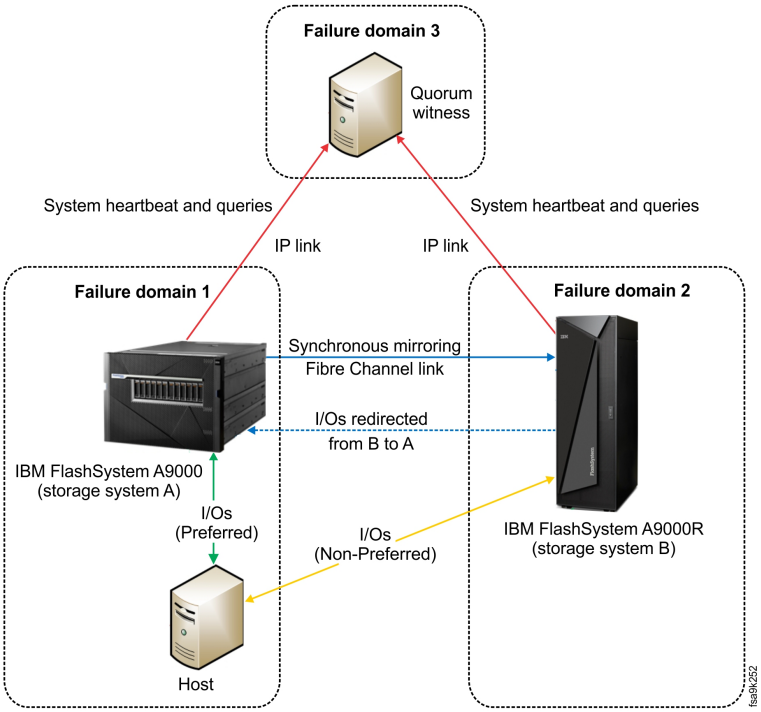

- **Quorum Witness and storage systems**
  A minimal IBM HyperSwap high availability configuration consists of two independently located IBM storage systems with configured synchronous replication accessible to the same hosts, and a Quorum Witness at a third site.
- **Multi-site HA/DR solution**
  The Multi-site function allows customers to deploy high availability (HA) and disaster recovery (DR) solutions over several sites, keeping multiple concurrent copies of data.

Quorum Witness and storage systems

A minimal IBM HyperSwap high availability configuration consists of two independently located IBM storage systems with configured synchronous replication accessible to the same hosts, and a Quorum Witness at a third site.

Important: Quorum Witness cannot reside on the same set of FlashSystem A9000 and/or A9000R systems of which it serves as part of a HyperSwap configuration.
After the Quorum Witness is installed, it is registered on both storage systems, participating in the HyperSwap relation. Each storage system can be configured with one Quorum Witness instance, identified by its unique ID.
Note: To ensure that failover is possible at any time, the Quorum Witness can be made more resilient by deploying it on a dedicated highly-available virtual machine. The system robustness can be achieved using VMware HA products, as described in Highly-available VM deployment.
After the Quorum Witness is configured and activated on the storage system, the storage system starts sending keepalive messages to be used by the Quorum Witness to monitor the array health. After the HyperSwap relation is established, it is validated by the storage systems, relying on the information provided by the Quorum Witness. Then each storage system verifies that the Quorum Witness reports the health status of its peer. This ensures that both storage systems are connected to the same Quorum Witness and operate properly.

Multi-site HA/DR solution

The Multi-site function allows customers to deploy high availability (HA) and disaster recovery (DR) solutions over several sites, keeping multiple concurrent copies of data.

A Multi-site HA/DR configuration is composed of three peers with one HyperSwap, employing a Quorum Witness and two asynchronous mirroring relations between them without participation of the Quorum Witness. One of these asynchronous mirrors remains on standby. This provides both failover and failback mechanisms, while keeping data mirrored to ensure business continuity, minimal risk for data outage, service availability, and data restoration in the event of disaster.
Important: Quorum Witness cannot reside on the same set of FlashSystem A9000 and/or A9000R systems of which it serves as part of a HyperSwap configuration.

Figure 1. Multi-site HA/DR diagram

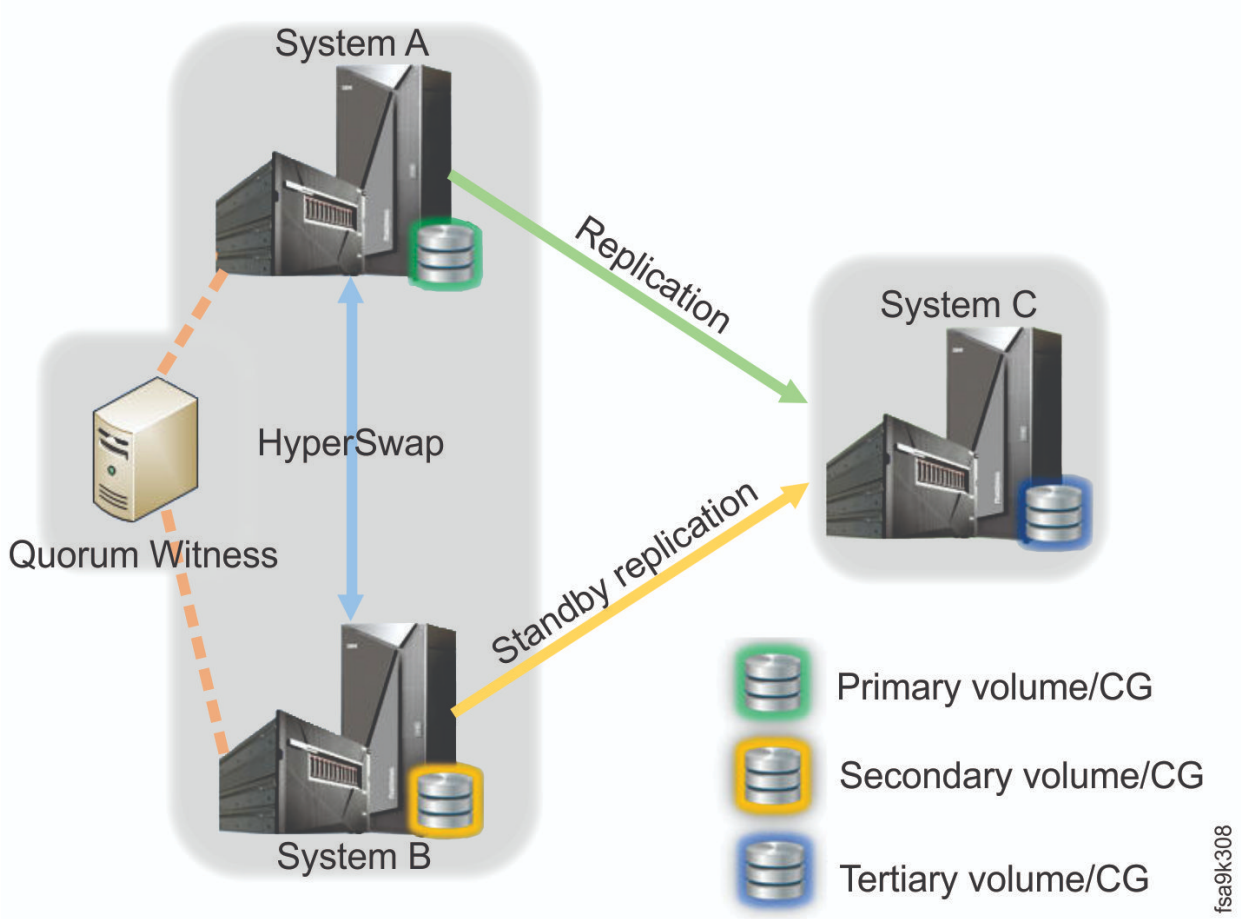

Multi-site HA/DR provides a higher level of data protection and more recovery capabilities than the common two-way high availability or mirroring solution. It fundamentally enhances business continuity, effectively prevents against extensive downtime, and allows to accommodate the most extensive compliance-related requirements.

Multiple Multi-site relations can exist on an IBM FlashSystem® A9000 or A9000R storage system, each with different peers. Also, an IBM FlashSystem A9000 or A9000R storage system can have different volumes with different roles in Multi-site HA/DR configurations.

Installation

Download and install the Quorum Witness software package as described in the following sections.

Important: Quorum Witness cannot reside on the same set of FlashSystem A9000 and/or A9000R systems of which it serves as part of a HyperSwap configuration. For information about uninstallation, see Uninstalling the Quorum Witness software.

- **Compatibility and requirements**
  For the complete and up-to-date information about the compatibility and requirements of IBM® Spectrum Accelerate Family HyperSwap Quorum Witness, refer to its latest release notes.
- **Performing first-time installation**
  You can install the Quorum Witness software on a bare-metal machine or a dedicated VM, by using a compatible version of Red Hat Enterprise Linux (RHEL) or CentOS.
- **Upgrading an existing installation**
  Use this information for upgrade to Quorum Witness 1.0.3.
- **Defining Quorum Witness on IBM storage systems**
  The Quorum Witness instance must be defined on the IBM FlashSystem A9000/A9000R storage systems, participating in a HyperSwap relation.
- **Uninstalling the Quorum Witness software**
  If you want to completely remove the Quorum Witness software from the Linux host upon which it is installed, follow the steps in the following procedure.

Compatibility and requirements

For the complete and up-to-date information about the compatibility and requirements of IBM® Spectrum Accelerate Family HyperSwap Quorum Witness, refer to its latest release notes.

Performing first-time installation

You can install the Quorum Witness software on a bare-metal machine or a dedicated VM, by using a compatible version of Red Hat Enterprise Linux® (RHEL) or CentOS.

Before you begin

Important: Quorum Witness cannot reside on the same set of FlashSystem A9000 and/or A9000R systems of which it serves as part of a HyperSwap configuration.

- Verify that the following packages are installed on the Quorum Witness host. These files are also supplied in the package together with the product BIN file. They can be deployed during the installation process, see step 5 of the following procedure.
  - pgdg-redhat-repo-latest.noarch.rpm
  - nginx-1.18.0-1.el8.ngx.x86_64.rpm
  - postgresql12-12.5-1PGDG.rhel8.x86_64.rpm
  - postgresql12-contrib-12.5-1PGDG.rhel8.x86_64.rpm
  - postgresql12-libs-12.5-1PGDG.rhel8.x86_64.rpm
  - postgresql12-server-12.5-1PGDG.rhel8.x86_64.rpm
  - uuid-1.6.2-42.el8.x86_64.rpm
  - jemalloc-5.2.1-2.el8.x86_64.rpm
  - redis-5.0.3-2.module+el8.0.0.z+3657+acb471dc.x86_64.rpm

  Note: In addition to the files above, be sure to also install the following:
  - The latest rsyslog RPM.
  - libxslt-1.1.32-5.el8.x86_64.
- A new Linux username – **ibmqw** – is created during installation to be used for the Quorum Witness management operations. You can customize the user ID for **ibmqw** by adding a Linux user (**useradd** command in RHEL) prior to the package installation. In this case, create the /home/ibmqw directory (or any other different path) before starting the installation process.
- Verify that the **ibmqw** user can access the /opt/ibm and /var/log/qw folders.

Procedure

Follow these steps to install the Quorum Witness.

1. Open the relevant TCP ports for a default Linux firewall:
   - 8460 for API communication
   - 8461 for the Quorum Witness log retrieval
   - 8462 for retrieval of Ngnix statistics (optional)

   Note: If you are using a different firewall software, refer to its documentation for the port opening instructions.
   - **firewall-cmd --permanent --add-port=8460/tcp**
   - **firewall-cmd --permanent --add-port=8461/tcp**
   - **firewall-cmd --permanent --add-port=8462/tcp**
   - **firewall-cmd --reload**
2. Download the installation package to a local folder on the Linux host that will be used as Quorum Witness server.
3. Extract the installation package file ('* ' represents the build number) :

   ```
   # tar -xzvf ibm_quorum_witness-1.0.3-*.x86_64.tar.gz
   ```

   The following files are extracted:
   - Dependencies (RPM files).
   - ibm_quorum_witness-1.0.3-*.x86_64.bin – product BIN file
   - ibm_quorum_witness-1.0.3-xxxx.x86_64.bin.asc– digital signature file for the BIN file verification
4. Verify that you have relevant permissions on the host.
5. If the complementary RPM files were not deployed before the installation of the Quorum Witness package, go to the extracted directory and use the **rpm -iv *.rpm** command to install them.
6. Enter **./ibm_quorum_witness-1.0.3-*.x86_64.bin** to start the installation.
7. Review and accept the license agreement which is displayed after you run the installation file.
   After installation, the Quorum Witness application is started automatically.
   Note: During the installation, a self-signed Secure Sockets Layer (SSL) certificate (**qw.crt**) and certificate key (**qw.key**) files are generated to create a secure communication channel for a storage system and the Quorum Witness host. Likewise, the Quorum Witness trusts the default certificate on the IBM storage system. If you need to replace any of these certificates, see the Managing SSL certificates section for instructions.
8. Copy the certificate file from the /opt/ibm/ibm_quorum_witness/settings/ssl_cert directory to each storage system, participating in the HyperSwap deployment. This certificate must be added to the storage system, when defining the Quorum Witness, using the command-line interface or IBM Hyper-Scale Manager.
9. In addition, store the **qw.crt** and **qw.key** files in a secure location, different from the Quorum Witness server or VM. These files may be required, when restoring the Quorum Witness database, if it becomes corrupted. Also, take note of the IP address or hostname, used during the installation. The same values must be used during software database restore as well. The Quorum Witness database restoration is detailed in Restoring Quorum Witness database.
10. If you are using SELinux, the program activation can fail. In this case allow nginx web server to bind to network interfaces and connect to the Quorum Witness socket.
    a. `semodule -i /opt/ibm/ibm_quorum_witness/conf.d/selinux/rhel8/ibmqw.pp`
    b. Verify the new configuration by viewing the ibmqw.te file. This file presents the changes to the relevant SELinux *.pp files in the human-readable format.
    `cat /opt/ibm/ibm_quorum_witness/conf.d/selinux/rhel8/ibmqw.te`
    c. Restart the Quorum Witness service, as explained in Checking and controlling the Quorum Witness service.

Upgrading an existing installation

Use this information for upgrade to Quorum Witness 1.0.3.

Quorum Witness 1.0.3 upgrades are supported with reinstallation.

If you are using a previous version of Quorum Witness, be sure to uninstall it before upgrading your operating system to RHEL or CentOS 8.2, using one of the following procedures.

Installing Quorum Witness 1.0.3 on a new system
Refer to Restoring the database by installing Quorum Witness on a new system.
Installing Quorum Witness 1.0.3 on an existing system
Refer to Restoring the database by reinstalling Quorum Witness on the existing system.
Important: During the reinstallation of Quorum Witness on the existing system, perform the RHEL/CentOS upgrade to version 8.2 after backing up the certificate and uninstalling the Quorum Witness.

Defining Quorum Witness on IBM storage systems

The Quorum Witness instance must be defined on the IBM FlashSystem A9000/A9000R storage systems, participating in a HyperSwap relation.

Important: Quorum Witness cannot reside on the same set of FlashSystem A9000 and/or A9000R systems of which it serves as part of a HyperSwap configuration. To ensure proper operation, define the Quorum Witness, using one of the following methods:

- IBM Hyper-Scale Manager, hybrid cloud management application for IBM Spectrum Accelerate Family products. Navigate to SYSTEMS & DOMAINS VIEWS > Quorum Witnesses and define Quorum Witness instances according to your application requirements.
  For the description of the IBM Hyper-Scale Manager operation, refer to its user guide.

  Figure 1. Navigating to Quorum Witness form in IBM Hyper-Scale Manager

  Figure 2. Quorum Witness form in IBM Hyper-Scale Manager

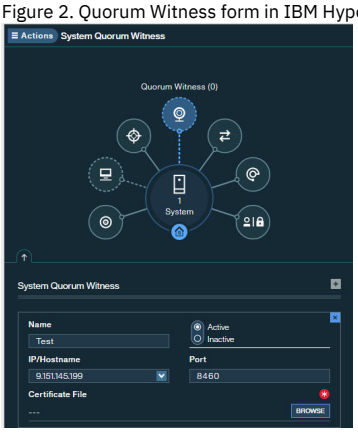

- XCLI utility, command-line user interface for managing IBM Spectrum Accelerate Family storage systems. Use the **quorum_witness_define** command for defining Quorum Witnesses.
  For details, refer to the IBM FlashSystem A9000/A9000R CLI reference guide.

Communication between the Quorum Witness and storage systems is performed over TCP/IP links. Verify that the TCP ports 8460 and 8461 are open on the external interface of the host. Refer to the installation procedure in see the Performing first-time installation section for the port opening procedures.

Uninstalling the Quorum Witness software

If you want to completely remove the Quorum Witness software from the Linux® host upon which it is installed, follow the steps in the following procedure.

Procedure

To uninstall the Quorum Witness software from the Linux host:

1. Log on to the Linux command prompt environment as a root user.
2. Run the standard Linux uninstallation command for the Quorum Witness software package (**rpm -e ibm_quorum_witness**).

Operation

This chapter describes operation tasks of managing SSL certificates and controlling the Quorum Witness service. For the detailed description of the CLI commands for managing and monitoring Quorum Witness from the storage system side, refer to the IBM FlashSystem A9000/A9000R CLI reference guide.

- **Managing SSL certificates**
  The Quorum Witness uses SSL certificates for maintaining a secure link to a storage system. This section describes how to manage the Quorum Witness SSL certificates.
- **Checking and controlling the Quorum Witness service**
  At any time, you can check whether the Quorum Witness service runs properly on the Linux® host. You can also stop and then restart the service, if needed.
- **Modifying the Quorum Witness communication port**
  At any time, you can modify the default communication port used by the Quorum Witness in accordance with your network requirements.
- **Enabling IPv6 connectivity**
  By default, the Quorum Witness uses the IPv4 communication method. If required, you can enable IPv6 connectivity.

Managing SSL certificates

The Quorum Witness uses SSL certificates for maintaining a secure link to a storage system. This section describes how to manage the Quorum Witness SSL certificates.

- **Replacing a self-signed certificate on the Quorum Witness**
  During the installation, a self-signed Secure Sockets Layer (SSL) certificate is generated to create a secure communication channel for a storage system and the Quorum Witness. If you already have a trusted certificate that you want to use, you can replace the self-signed certificate with an existing certificate on the Quorum Witness.
- **Managing CA public key on the Quorum Witness**
  If a CA public key certificate is replaced on a storage system, it must be added to the trust store on the Quorum Witness to ensure a trusted link between the Quorum Witness and the storage system.

Replacing a self-signed certificate on the Quorum Witness

During the installation, a self-signed Secure Sockets Layer (SSL) certificate is generated to create a secure communication channel for a storage system and the Quorum Witness. If you already have a trusted certificate that you want to use, you can replace the self-signed certificate with an existing certificate on the Quorum Witness.

Procedure

A self-signed certificate file, **qw.crt**, and a certificate key file, **qw.key**, are stored in the following directory: /opt/ibm/ibm_quorum_witness/settings/ssl_cert. The following procedure describes how to replace them.
Important: When there are existing HyperSwap relations on the systems registered on the Quorum Witness, they must be deactivated before and reactivated after the certificate change. For more information, see the known issues of the release notes.

1. Create new X.509-type certificate and key files, using a utility of your choice. Make sure to use the same names for both files. Alternatively, you can use existing files in the PEM format.
2. Copy the new certificate and key files to the /opt/ibm/ibm_quorum_witness/settings/ssl_cert directory, overwriting the old copies.
3. Restart the nginx service, using the **service nginx restart** command.
4. Add the public key of the new certificate file to each storage system, participating in the HyperSwap deployment. This certificate must be added to the storage system, when defining the Quorum Witness. This procedure is explained in the IBM FlashSystem A9000/A9000R CLI reference guide or IBM Hyper-Scale Manager user guide, depending on the management method that you intend to use.

Managing CA public key on the Quorum Witness

If a CA public key certificate is replaced on a storage system, it must be added to the trust store on the Quorum Witness to ensure a trusted link between the Quorum Witness and the storage system.

Procedure

The IBM CA public key, **ibm-xiv.crt**, is stored in the following directory: /opt/ibm/ibm_quorum_witness/settings/ssl_cert/client_certificates. The following procedure describes how to remove the old key, add the new one and list all keys stored on the Quorum Witness.

1. Remove the old key by running the following command: **su - ibmqw -c "qw_certificates remove -f ibm-xiv.crt"**.
2. Copy the new CA public key to any location on the Quorum Witness host and run the following command to add the key: **su - ibmqw -c "qw_certificates add -f <path_to_new_certificate>/<new_certificate_name>"**.
   The new CA public key is added to the database and copied to the /opt/ibm/ibm_quorum_witness/settings/ssl_cert/client_certificates directory.
3. Restart the nginx service, using the **service nginx restart** command.
4. Display the trusted certificate to verify that the operation succeeded. Run this command: **su - ibmqw "qw_certificates list"**. The output must be as follows:

```
ibmqw@test-qw2 ~]$ qw_certificates list
qw.crt
ibm-xiv.crt
```

Checking and controlling the Quorum Witness service

At any time, you can check whether the Quorum Witness service runs properly on the Linux® host. You can also stop and then restart the service, if needed.

Procedure

1. Log on to the Linux command prompt environment as a root user.
   Important: Only root users can complete service operations.
2. Enter the following command: **ibm_quorum_witness status**.
   The status of the service is displayed.

```
# service ibm_quorum_witness status
Celery services are running...
Django service is running...
```

What to do next

If you want to stop the Quorum Witness service, use the **stop** command:
Important: Stopping the service makes the Quorum Witness unavailable for storage systems, disrupting the HyperSwap operation.

```
# service ibm_quorum_witness stop
```

To start the service again, use the **start** command:

```
# service ibm_quorum_witness start
```

If you want to stop and then restart the Quorum Witness service in one command, use the **restart** command:

```
# service ibm_quorum_witness restart
```

Modifying the Quorum Witness communication port

At any time, you can modify the default communication port used by the Quorum Witness in accordance with your network requirements.

Procedure

1. Log on to the Linux® command prompt environment as a root user.
2. Enter the following command: **sed --in-place --copy 's/current_port/new_port/g' /opt/ibm/ibm_quorum_witness/conf.d/nginx/qw_nginx.conf**.
3. Use the same command to change the current port incremented by 1 to the new port incremented by 1. This port is reserved for log retrieval by storage systems in the HyperSwap solution.
4. Restart the nginx service, using the **service nginx restart** command.
   For example, to change port 8460 to port 8470:

```
sed --in-place --copy 's/8460/8470/g' /opt/ibm/ibm_quorum_witness/conf.d/nginx/qw_nginx.conf
sed --in-place --copy 's/8461/8471/g' /opt/ibm/ibm_quorum_witness/conf.d/nginx/qw_nginx.conf
service nginx restart
```

5. Open the new ports, as explained in the Performing first-time installation section.

Enabling IPv6 connectivity

By default, the Quorum Witness uses the IPv4 communication method. If required, you can enable IPv6 connectivity.

Procedure

1. Log on to the Linux® command prompt environment as a root user.

2. Enter the following command: **sed --in-place --copy 's/#listen \[::\]/listen \[::\]/g' /opt/ibm/ibm_quorum_witness/conf.d/nginx/qw_nginx.conf**.
3. Restart the nginx service, using the **service nginx restart** command.

Troubleshooting

This chapter can help you detect and solve problems that you might encounter when using the IBM® Spectrum Accelerate Family HyperSwap Quorum Witness.

Note:

- For up-to-date information about known issues and possible workarounds, refer to the latest Release Notes.
- When contacting IBM Support, specify the storage system you are managing, using Quorum Witness.

See the following sections for more information:

- **Checking the log files**
  The IBM Spectrum Accelerate Family HyperSwap Quorum Witness maintains log files that record different types of events.

Checking the log files

The IBM® Spectrum Accelerate Family HyperSwap Quorum Witness maintains log files that record different types of events.

You can find the following log files in the /var/log/qw/ directory:

- events.log – Records all events according to their type: **Info**, **Error**, or **Warning**. You can use the Rsyslog application, an open source utility for forwarding log messages over IP networks (for more information, see the Rsyslog website). The following example shows different events that might be recorded:

```
Mar 14 14:44:41 qw-rh82-1 INFO 52 | external | quorum_witness | Service started
Mar 15 10:12:53 qw-rh82-1 INFO 53 | external | quorum_witness | Service stopped
Mar 19 08:30:39 qw-rh82-1 INFO 54 | external | 7023804-2 | System 7023804-2 registered
Mar 19 15:13:59 qw-rh82-1 INFO 55 | external | 7023804-2 | System 7023804-2 unregistered
```

- qw.log – Records events related to the operation of the Quorum Witness application.
- celery.log – Records events regarding the Celery services on the Linux® host on which Quorum Witness is installed.
- django.log – Records events regarding the Django service on the Linux host on which Quorum Witness is installed.
- celerybeat.log – Records events regarding the scheduled tasks of the Celery services on the Linux host.
- gunicorn.log – Records application-level events related to the Django service on the Linux host.
- nginx_access.log – Records access events of the NGINX service on the Linux host.
- nginx_error.log – Records errors of the NGINX service on the Linux host.
- redis.log – Records Redis-related events on the Linux host.

Best practices

Refer to the general guidance and best practices that are described in the following sections.

- **Highly-available VM deployment**
  You can utilize vSphere Fault Tolerance (FT) for the Quorum Witness virtual machine to improve the HyperSwap application reliability and decrease the system downtime.
- **Restoring Quorum Witness database**
  The Quorum Witness database may become corrupted, making it unavailable for the storage systems and disrupting the HyperSwap operation. Database corruption can occur due to the physical failure of the Quorum Witness server hard disk. Also, it may happen during incorrect execution of VM failover, while using vSphere Fault Tolerance protection technique. The database failure is indicated by the DB Recovery Needed status of the Quorum Witness in the IBM Hyper-Scale Manager GUI or XCLI and relevant information message stored in the log file.
- **Additional reference information**
  IBM provides additional reference documentation to help users integrate, implement and operate realistic scenarios for IBM storage systems in a HyperSwap solution.

Highly-available VM deployment

You can utilize vSphere Fault Tolerance (FT) for the Quorum Witness virtual machine to improve the HyperSwap application reliability and decrease the system downtime.

FT provides continuous availability for the Quorum Witness VM by creating and maintaining another VM that is identical and continuously available to replace it in the event of a failover situation. The duplicate virtual machine is created and runs on another host. Its execution is identical to that of the primary VM and it can take over at any point without interruption, thereby providing fault tolerant protection.

Important: Quorum Witness cannot reside on the same set of FlashSystem A9000 and/or A9000R systems of which it serves as part of a HyperSwap configuration. The Fault Tolerance is a preferred solution for the Quorum Witness VM protection. It ensures zero downtime and no impact on the operation in case of the Quorum Witness host and/or datastore migration.

Restoring Quorum Witness database

The Quorum Witness database may become corrupted, making it unavailable for the storage systems and disrupting the HyperSwap operation. Database corruption can occur due to the physical failure of the Quorum Witness server hard disk. Also, it may happen during incorrect execution of VM failover, while using vSphere Fault Tolerance protection technique. The database failure is indicated by the DB Recovery Needed status of the Quorum Witness in the IBM Hyper-Scale Manager GUI or XCLI and relevant information message stored in the log file.

Before you begin

Important: Quorum Witness cannot reside on the same set of FlashSystem A9000 and/or A9000R systems of which it serves as part of a HyperSwap configuration.

Procedure

To restore the corrupted Quorum Witness database:

1. Reinstall the Quorum Witness software on a bare-metal machine or a dedicated VM, using the same IP address or hostname as the previous installation. See Performing first-time installation.
2. Copy the previously saved qw.cert and qw.key files to the/opt/ibm/ibm_quorum_witness/settings/ssl_cert directory.
3. Restart the nginx service, using the **service nginx restart** command.
4. Activate the Quorum Witness via XCLI or Hyper-Scale Manager.

- **Restoring the database by installing Quorum Witness on a new system**
  Complete the following steps to restore the Quorum Witness database by installing Quorum Witness on a new system.
- **Restoring the database by reinstalling Quorum Witness on the existing system**
  Complete the following steps to restore the Quorum Witness database by reinstalling Quorum Witness on the existing system.

Restoring the database by installing Quorum Witness on a new system

Complete the following steps to restore the Quorum Witness database by installing Quorum Witness on a new system.

Before you begin

Important: Quorum Witness cannot reside on the same set of FlashSystem A9000 and/or A9000R systems of which it serves as part of a HyperSwap configuration.

Procedure

To restore the corrupted Quorum Witness database by reinstalling Quorum Witness on a new system:

1. Reinstall the Quorum Witness software on a bare-metal machine or a dedicated VM, using the same IP address or hostname as the previous installation. See Performing first-time installation.
2. Copy the previously saved qw.cert and qw.key files to the/opt/ibm/ibm_quorum_witness/settings/ssl_cert directory.
3. Restart the nginx service, using the **service nginx restart** command.
4. Activate the Quorum Witness via XCLI or Hyper-Scale Manager.

Restoring the database by reinstalling Quorum Witness on the existing system

Complete the following steps to restore the Quorum Witness database by reinstalling Quorum Witness on the existing system.

Procedure

To restore the corrupted Quorum Witness database by reinstalling Quorum Witness on the existing system:

1. Back up the certificate directory using the following command:
   ```
   cp -r
   /opt/ibm/ibm_quorum_witness/settings/ssl_cert /tmp
   ```
2. Verify that you have the installation file with the same level by running the following command:
   ```
   rpm -qa | grep quorum
   ```
3. Uninstall the existing Quorum Witness by running the following command:
   ```
   rpm -e ibm_quorum_witness
   ```
4. Reinstall Quorum Witness by using the steps provided in Performing first-time installation.
5. Stop the Quorum Witness service by running the following command:
   ```
   systemctl stop ibm_quorum_witness
   ```
6. Replace the old qw.cert and qw.key files by running the following commands:

   ```
   cp /tmp/qw.key /opt/ibm/ibm_quorum_witness/settings/ssl_cert/qw.key
   cp /tmp/qw.crt /opt/ibm/ibm_quorum_witness/settings/ssl_cert/qw.crt
   ```

7. Restart the Quorum Witness service by running the following command:
   `systemctl start ibm_quorum_witness`

Additional reference information

IBM provides additional reference documentation to help users integrate, implement and operate realistic scenarios for IBM storage systems in a HyperSwap solution.

IBM Redbooks are technical documents where IBM® experts share their expertise and best practices for using IBM Spectrum Accelerate Family HyperSwap Quorum Witness, as an active element of the IBM HyperSwap solution. Among others, you can refer to this publication – IBM HyperSwap for IBM FlashSystem A9000 and A9000R.

Operating

Use these operating procedures to start and stop the system.

- **Powering on after shutdown**
  Use this procedure to power on a FlashSystem A9000R system after either someone shut down the system or after the system automatically shut down.
- **Powering off the system**
  Use this procedure to power off a FlashSystem A9000R system by initiating a controlled shutdown.

Powering on after shutdown

Use this procedure to power on a FlashSystem A9000R system after either someone shut down the system or after the system automatically shut down.

About this task

IBM FlashSystem® A9000R automatically performs an emergency shutdown when there is a power failure, when various systems and/or components have failed, or when a system thermal threshold has been exceeded (above allowed temperature). Emergency shutdowns are not performed by personnel. The system determines when to carry out an emergency shutdown.

If an external power loss occurs, the system will go into emergency shutdown. This will occur if the emergency shutdown was completed successfully (no interrupts of any kind). Should the system exhibit issues during the power-on operation (partial power on, boot to maintenance state, etc.), avoid all manual intervention and contact the next level of support.
Note: If the system loses external power, the system will automatically shut down. When the external power returns, the system will automatically power up.
Attention: If you are recovering an IBM FlashSystem A9000R storage system after an emergency shutdown, determine the reasons for the shutdown. Correct the cause for the emergency shutdown before you power on the system. If you cannot determine why the emergency shutdown occurred, contact an IBM® service representative for assistanceyour next level of support.

Procedure

Complete the following steps to power on the system:

Note: Depending on the rack configuration, different racks have different front panels.

- The perforated panels allow PDU LED indications to be seen through the panel, while the louvered panels often need to be removed.
- Louvered panels should be removed at the beginning of this procedure for easier viewing.

1. If necessary, remove the louvered panel from the front of the system by using a Phillips-head screwdriver to loosen the captured screws.
2. Ensure that external electrical sources are correctly supplying power to the storage system.
3. Verify that the customer-site circuit breakers are on.
   If the circuit breakers are not on, contact your IBM service representativeyour next level of support.

4. Verify that the power cables are connected correctly.
5. Verify that LEDs on the front of the PDU are lit, indicating that the PDU is receiving AC power, see 6 Figure 1

Figure 1. Front of the PDU

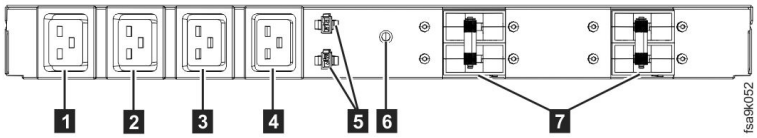

6. Flip the system ON/OFF button to the **O** (off) position (see Figure 2).
   Wait at least 5 seconds.

Figure 2. Main circuit breaker on front of the rack

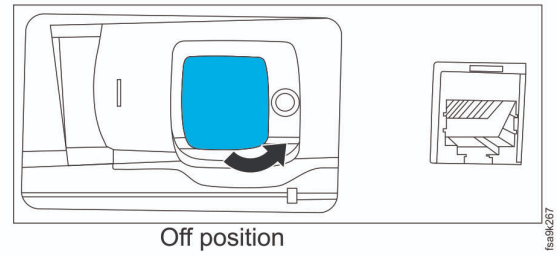

7. Turn on the system by flipping the switch to the **|** (on) position.
8. Place the louvered panel back on the front of the system by using a Phillips-head screwdriver.

Powering off the system

Use this procedure to power off a FlashSystem A9000R system by initiating a controlled shutdown.

Before you begin

Attention: Power off the FlashSystem A9000R system only when necessary. When powered off, the storage system is taken out of use and can no longer store data from host systems. Taking the system out of use might put extra burden on all other FlashSystem A9000R systems and might affect storage performance for the site.
Important:

- Although the IBM® MicroLatency modules retain data if the enclosure is temporarily disconnected from power, if the system is powered off for 90 days, data errors might occur. The storage enclosure can be safely powered down for up to 90 days, in temperatures up to 40°C (104°F). When the enclosure is first installed at the customer site, it also automatically reformats all the flash memory, providing improved customer experience.
- If the system is powered off for long periods of time (longer than 9 months), the grid controller battery modules may become unusable.
  Help prevent this by making sure that the batteries are fully charged before powering off the system. If the batteries did not fully charge within 2 hours after power up, replace, as needed.

  For systems containing battery module part numbers 00AR056 or 0AR260: If the system is powered off for a period that exceeds 40 days, the grid controller battery modules may become unusable. To prevent this, the system must be connected to power or powered on for two hours (at least) within each storage month.

Note: If the storage enclosure has been powered off for longer than seven days, powering on the system automatically starts a proprietary deep scrub and refresh process. This action restores the system to optimal condition without customer intervention. The storage enclosure also employs several powerful methods for recovering data in situations where the system is stored for periods longer than 90 days, or at temperatures higher than 40°C. Every possible effort is made to recover any lost data.

About this task

The storage system automatically performs an emergency shutdown when there is a power failure or when when various systems and/or components have failed. An emergency shutdown is not performed by personnel. The system determines when to carry out an emergency shutdown.

When you power off the storage system in a controlled fashion, the FlashSystem A9000R completes the following actions:

- Stops serving hosts
- Flushes all data from the cache, and places all the data onto the SSDs (also called *vaulting data*)
- Powers down the flash enclosures

For CLI instructions for shutting down the system, see "Shutting down the system" in the CLI reference information.

Managing

Refer to the following sections for information about managing IBM FlashSystem A9000R systems using IBM Hyper-Scale Manager and the command-line interface (CLI).

- **Managing systems using IBM Hyper-Scale Manager**
  Use IBM® Hyper-Scale Manager to manage IBM Spectrum Accelerate systems.
- **Managing systems using the command-line interface (CLI)**
  For information about how to manage IBM FlashSystem® A9000R systems using the command-line interface (CLI), refer to the appropriate following section.

Managing systems using IBM Hyper-Scale Manager

Use IBM® Hyper-Scale Manager to manage IBM Spectrum Accelerate systems.

For full product information, see IBM Hyper-Scale Manager documentation.

Managing systems using the command-line interface (CLI)

For information about how to manage IBM FlashSystem® A9000R systems using the command-line interface (CLI), refer to the appropriate following section.

For 12.3.2.x command-line line information, see Command-line reference (version 12.3.2.x), in Reference.
Note: Download older command-line information from Publications and related information.

Configuring

Configure host connections and key servers on your FlashSystem A9000R system.

Configuring host attachment

Configure various parameters for host attachment.

Hosts can be attached by Fibre Channel and Ethernet iSCSI. Refer to the subtopics of this section for information about setting host attachment parameters for various operating systems.
Important: Always use your operating system host connectivity options together with the multipath best practices.
For more information about physical connectivity, zoning, and best practices, see Host network connectivity and multipathing configurations.

Direct connectivity support

The operating systems that support direct connectivity include, *but are not limited*, to the following:
Note: For a full list of the latest supported operating systems, see System Storage® Interoperation Center website .

Linux®

- RHEL 8.4 — FC and Ethernet

Microsoft Windows

- Windows 2019 — FC and Ethernet
- Windows 2022 — FC and Ethernet

Solaris

- Solaris 10 — FC
- Solaris 11 — FC

- **Preparing for host attachment**
  Preparation is may be necessary before host attachment to your storage system.
- **Defining the host on the storage system**
  Manually define the host on the storage system using either the command-line interface (CLI) or Hyper-Scale Manager. After the host is defined on the storage system, map storage volumes to the host using the Hyper-Scale Manager.
- **Monitoring host connectivity**
  Monitor the connectivity of attached hosts in real-time through either the Hyper-Scale Manager or the command-line interface (CLI).
- **Configuring Linux host attachment**
  Configure a host server using a Linux operating system for host attachment to your storage system.
- **Configuring Windows host attachment**
  Configure a host server using a Windows operating system for host attachment to your storage system.
- **Configuring Solaris host attachment**
  Connect your storage system to a Solaris 10 or 11 Fibre Channel (FC) host operating system.

Preparing for host attachment

Preparation is may be necessary before host attachment to your storage system.

The preparation stages are:

1. Physically connecting the host to the storage system
2. Host connectivity and parameter settings for FlashSystem A9000 and A9000R

Note: Perform these preparation stages only once per host or cluster. If you have already performed these preparation stages, skip to the host attachment instructions for your operating system.

- **Physically connecting the host to the storage system**
  Physically connect hosts to the storage system before operating system host attachment.
- **Host connectivity and parameter settings for FlashSystem A9000 and A9000R**
  After you physically connect your host or cluster to the IBM storage system, you must ensure that the connection type (FC or iSCSI) is properly defined on

that system. This section details general rules and recommendations for host connectivity and resiliency for the IBM FlashSystem A9000 and A9000R storage systems.

Physically connecting the host to the storage system

Physically connect hosts to the storage system before operating system host attachment.

Attention: Although you can use both FC and iSCSI to connect different hosts to the same IBM® storage system that supports both connection methods, use only one connection type per single host or cluster to avoid possible issues.

- **Physically connecting a host using Fibre Channel (FC) connections**
  The Fibre Channel (FC) protocol (FCP) is used to establish connectivity over FC infrastructure.
- **Physically connecting a host using iSCSI connections**
  The iSCSI protocol is used to establish connectivity over Ethernet infrastructure.

Physically connecting a host using Fibre Channel (FC) connections

The Fibre Channel (FC) protocol (FCP) is used to establish connectivity over FC infrastructure.

Hosts can be connected to the storage system via a private FC storage fabric, through an FC switch (not via a direct connection).

For the physical connection on the host side, install at least one host bus adapter (HBA) on the host. The HBA provides physical FC ports to which you can connect FC cables. FlashSystem A9000R supports different HBA brands and types. For the latest support information and compatibility matrix, refer to the System Storage® Interoperation Center website .
Note: For best performance, install the latest firmware and drivers for the HBAs that are in use. The HBA vendor should provide the latest firmware and drivers.
Important: A high availability (HA) configuration is the best practice for FC connectivity. For best performance, use a dedicated FC connectivity infrastructure.
For information on how to physically connect your host(s) to the supported storage system over Fibre Channel, see the following:

- Network and host connectivity requirements
- IBM® FlashSystem A9000, IBM FlashSystem® A9000R, and IBM XIV® Storage System: Host Attachment and Interoperability Redbook, SG24-8368

Physically connecting a host using iSCSI connections

The iSCSI protocol is used to establish connectivity over Ethernet infrastructure.

Hosts can be connected to the storage system via a private Ethernet-based LAN or WAN, through an Ethernet switch (not via a direct connection).
Important: A high availability (HA) configuration is the best practice for iSCSI connectivity. For best performance, use a dedicated iSCSI connectivity infrastructure.
For information on how to physically connect your host(s) to the supported storage system over Ethernet-based networks, see the following:

- Network and host connectivity requirements
- IBM® FlashSystem A9000, IBM FlashSystem® A9000R, and IBM XIV® Storage System: Host Attachment and Interoperability Redbook, SG24-8368

Host connectivity and parameter settings for FlashSystem A9000 and A9000R

After you physically connect your host or cluster to the IBM® storage system, you must ensure that the connection type (FC or iSCSI) is properly defined on that system. This section details general rules and recommendations for host connectivity and resiliency for the IBM FlashSystem® A9000 and A9000R storage systems.

IBM FlashSystem A9000 and A9000R storage systems communicate with their hosts over a Fibre Channel (FC) network or over an Ethernet network, using iSCSI protocol. Network connectivity is supported via a switching fabric, without direct host attachment. Host traffic can be directed to any of the grid controllers.

For full information on physical connectivity between storage systems and hosts, see Host network connectivity and multipathing configurations.

Achieving optimal multipathing

The administrator must ensure that multipathing is used. Successful multipathing is achieved by optimal host connectivity, adequate balancing of host workloads and other means, as detailed below.

Physical connectivity

When planning physical connectivity between storage systems and hosts, spread grid controller port connections as evenly as possible throughout the system and across the hosts. To achieve this, when connecting hosts to the storage system, the storage administrator must guarantee:

- Host connectivity without a single point of failure.
- Host workloads that are adequately balanced across the connections and grid controllers.
- Optimal way of using grid controller resources. This must be combined with periodical review of load balancing, especially when connections or traffic patterns change.

- Number of paths, in accordance with the minimum and recommended requirements.

Logical connectivity

Follow these rules when planning logical connectivity between storage systems and hosts.

- To avoid host performance degradation, do not exceed the total of six paths per host for FlashSystem A9000 and eight paths for FlashSystem A9000R.
- For better performance, balance your traffic across several storage system grid controllers.
- Configure at least one FC zone per host and per initiator.

The load balancing is achieved by using multipathing. In a production environment, it is important to connect hosts to a minimum of two separate SAN switches in independent fabrics to ensure redundancy.

Port combinations

IBM FlashSystem A9000 and A9000R storage systems have the following port combinations:

- A9000 – 12 FC or iSCSI ports.
- A9000R – 4 × number of grid controllers in the rack.

Each module has two dual-port adapters. The first adapter hosts ports 1 and 2. The second adapter hosts ports 3 and 4. By default, FC ports 1, 2 and 4 are designated as FC targets, and FC port 3, as initiator. You can change the roles, using CLI commands or the Hyper-Scale Manager.
See the following sections for more details regarding Fibre Channel connectivity information and configuration:

- Fibre Channel host port configuration
- Fibre Channel connectivity requirements

See the following sections for more details regarding iSCSI connectivity information and configuration:

- iSCSI host port configuration
- iSCSI connectivity requirements

Minimum connectivity requirements

The minimum number of paths required for system resiliency depends on the storage system type:

- FlashSystem A9000: two, one path to each switch fabric
- FlashSystem A9000R: four, two paths to each switch fabric

Important: Host multipath connectivity eliminates the risk of a single point-of-failure between the host and storage systems.

- **Setting Fibre Channel connectivity**
  Fibre Channel ports are identified by their World Wide Port Name (WWPN). The FC ports on the storage system must be identified before setting FC zoning on your hosts and the storage system.
- **Setting iSCSI connectivity**
  The FlashSystem A9000R storage system acts as a Transmission Control Protocol (TCP) server for iSCSI connections.

Setting Fibre Channel connectivity

Fibre Channel ports are identified by their World Wide Port Name (WWPN). The FC ports on the storage system must be identified before setting FC zoning on your hosts and the storage system.

Identify the FC ports from the IBM® Hyper-Scale Manager application:

1. Log in to the IBM Hyper-Scale Manager application and select a storage system that you want to configure.
2. From the System mini-dashboard, click the System Ports relation icon to display the list of FC ports on the system patch panel.

Figure 1. System mini-dashboard

The list of FC ports is displayed.

3. Click the Actions button on any particular FC port and select View/Edit FC Port to reveal the port parameters, including the WWPN.

Figure 2. Viewing FC port parameters

The FC port parameters are displayed.

4. In the FC Port form, identify the WWPN of the port.

Figure 3. FC port parameters

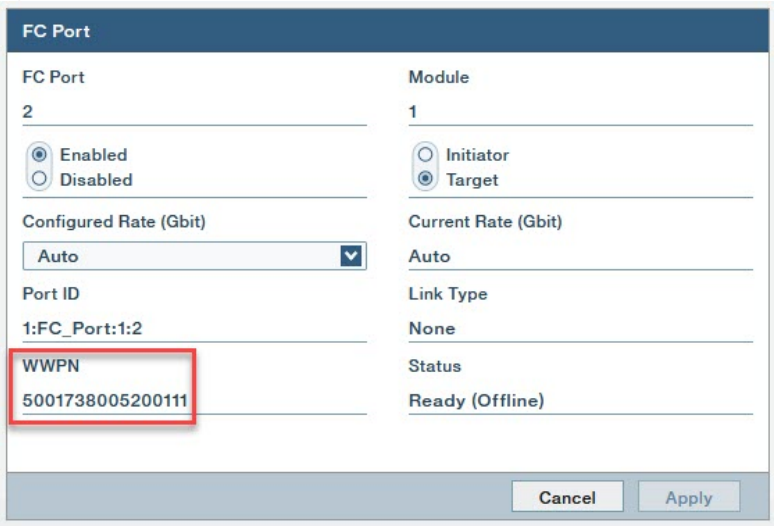

- **Fibre Channel zoning**
  Zoning is a connection isolation technique that is required when connecting hosts with the storage system over Fibre Channel SAN.

Fibre Channel zoning

Zoning is a connection isolation technique that is required when connecting hosts with the storage system over Fibre Channel SAN.

Zoning achieves the following:

- Creates a barrier between different network environments. Only the members of the same zone can communicate within that zone, and all external communications are blocked.
- Isolates any single host bus adapter (HBA) for security and reliability reasons.
- Allows finer segmentation of the switched fabric.

With proper zoning, you can avoid the following situations:

- HBAs of different vendors operate differently upon error recovery. This may impact other hosts connected to the same switch if those are not isolated through zoning.
- Any change in the SAN fabric, such as a change caused by a server restart or a new product being added to the SAN – triggers a registered state change notification (RSCN). An RSCN requires any change-aware device to acknowledge the change while interrupting its own traffic flow.

The most secure zoning is to have zones that consist of a single initiator and multiple targets.

For more information about zoning, see:

- Fibre Channel zone types
- IBM® FlashSystem A9000, IBM FlashSystem® A9000R, and IBM XIV® Storage System: Host Attachment and Interoperability Redbook, SG24-8368

General zoning guidelines

To achieve load balancing, the hosts must spread their connections among all grid controllers. This means that each host is zoned to a subset of targets on a certain combination of grid controllers in such a way, that all grid controllers have the same amount of incoming host paths. This can be gained by using *zone types*.

Each *zone type* is the set of targets (on certain combination of grid controllers) to which a host is zoned. In this manner, when zoning a host to the storage system, you must create a new zone, containing the host initiator and all the targets that reside in a specific *zone type*, to which the host is assigned.

The amount of *zone types* depends on the amount of paths being used for connectivity.

Setting iSCSI connectivity

The FlashSystem A9000R storage system acts as a Transmission Control Protocol (TCP) server for iSCSI connections.

All iSCSI packets are routed through the Ethernet port from which the iSCSI connection was initiated.

The following table specifies the primary settings that must be configured on the storage system for proper iSCSI connectivity with the host.

Table 1. Primary settings for proper iSCSI connectivity

Setting type	Configuration on the storage system
iSCSI qualified name (IQN)	Required for each host that uses iSCSI to connect to the storage system. The IQN uniquely identifies the different entities that utilize iSCSI on the IP network, and enables the operation of the iSCSI initiator. The IQN of the storage system is a pre-configured factory setting that cannot be changed.
IP interface	iSCSI uses IP-based communication. On the storage system, each iSCSI port is defined as an IP interface with its own IP address. Alternatively, the ports can be bundled (bonded) for load balancing to a single logical iSCSI port with a single IP address. By default, six iSCSI target ports are predefined on the storage system to serve hosts over iSCSI.
Default gateway IP address	Required if the connected hosts do not reside on the same Layer-2 subnet.
Maximum transmission unit (MTU)	Required if your network supports an MTU that is larger than the standard one. The MTU default value is 1,536 bytes and the maximum value is 8,192 bytes. You should specify the largest possible MTU. If supported by the switches and routers, use up to 8,192 bytes.

- **Configuring iSCSI ports**
  Perform the following procedure to configure iSCSI ports from the IBM Hyper-Scale Manager application.

Configuring iSCSI ports

Perform the following procedure to configure iSCSI ports from the IBM® Hyper-Scale Manager application.

Procedure

1. Log in to the IBM Hyper-Scale Manager application and select a storage system that you want to configure.
2. From the System mini-dashboard, click the System Ports relation icon to display the list of iSCSI ports on the system patch panel.

Figure 1. System mini-dashboard

The list of iSCSI ports is displayed.

3. Click the Actions button an iSCSI port and select Add IP Interface to display the port parameters and update the IP interface.

Figure 2. Displaying iSCSI port parameters

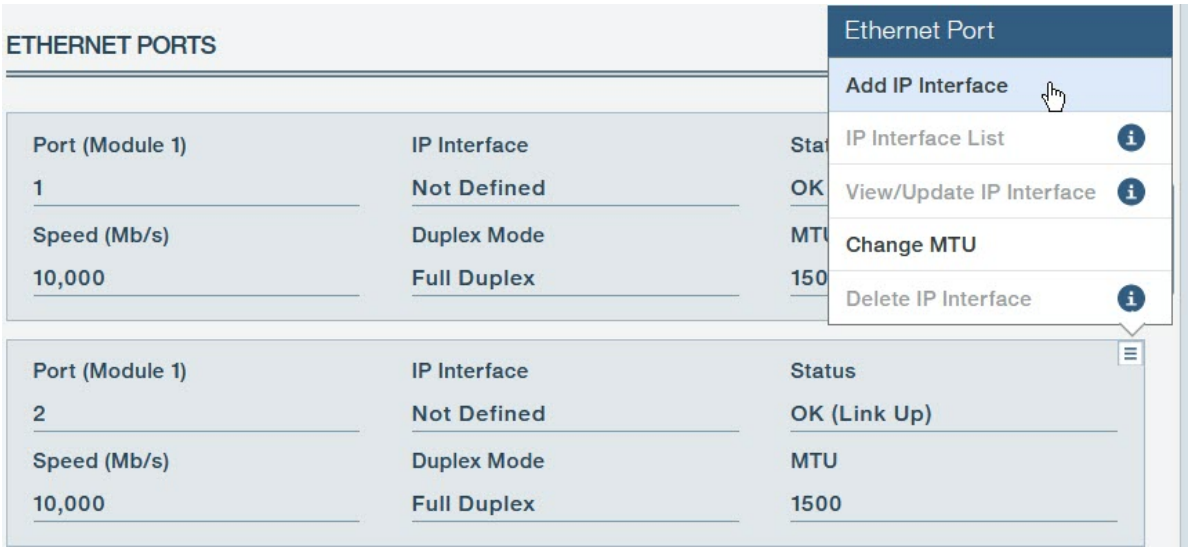

The iSCSI port parameters are displayed.

4. Fill in the Add IP Interface form.
   Use the information in Table 1 to fill in the following fields.
   - IP Interface Name – Define the name of the iSCSI interface.
   - VLAN Tag – Select if the VLAN is tagged or untagged.
   - IPv4 Address – IP address of the iSCSI interface.
   - Default Gateway – If needed, enter the IP address of the default gateway.
   - Netmask – Enter the netmask address, or select an existing one from the list box.

   Figure 3. Adding an IP interface for iSCSI port parameters

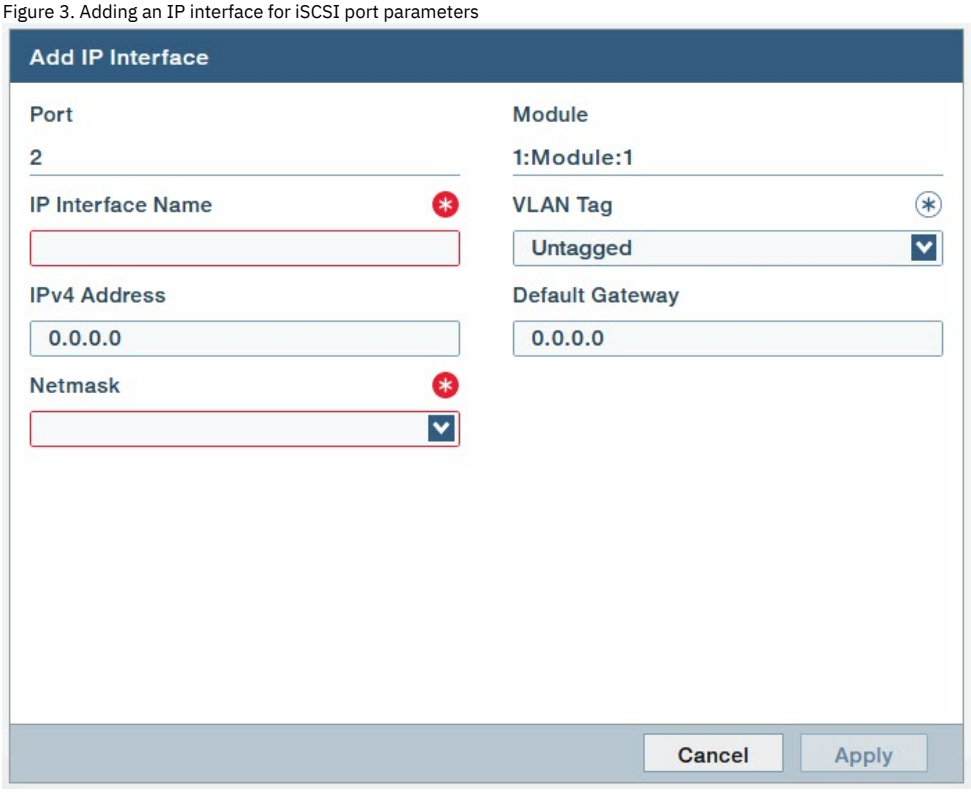

   Once set, click Apply.

5. Check and update, as needed, the MTU settings.
   Click the Actions button an iSCSI port and select Change MTU.
   Note: All devices in a network must use the same MTU. If in doubt, set MTU to 1500, because 1500 is the default value for Gigabit Ethernet. Performance might be impacted if the MTU is set incorrectly.
   Click Apply, if changes were made or click Cancel to exit without making any changes.

Defining the host on the storage system

Manually define the host on the storage system using either the command-line interface (CLI) or Hyper-Scale Manager. After the host is defined on the storage system, map storage volumes to the host using the Hyper-Scale Manager.

Defining a host from the Hyper-Scale Manager

1. Log in to the IBM® Hyper-Scale Manager application.
2. Click the New icon on the toolbar at the top and then select Host.

Figure 1. New Host option

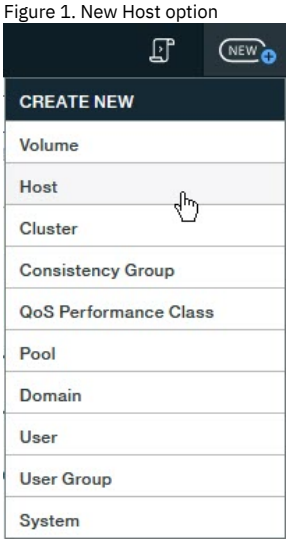

The Add Host form is displayed.

Figure 2. Add Host form

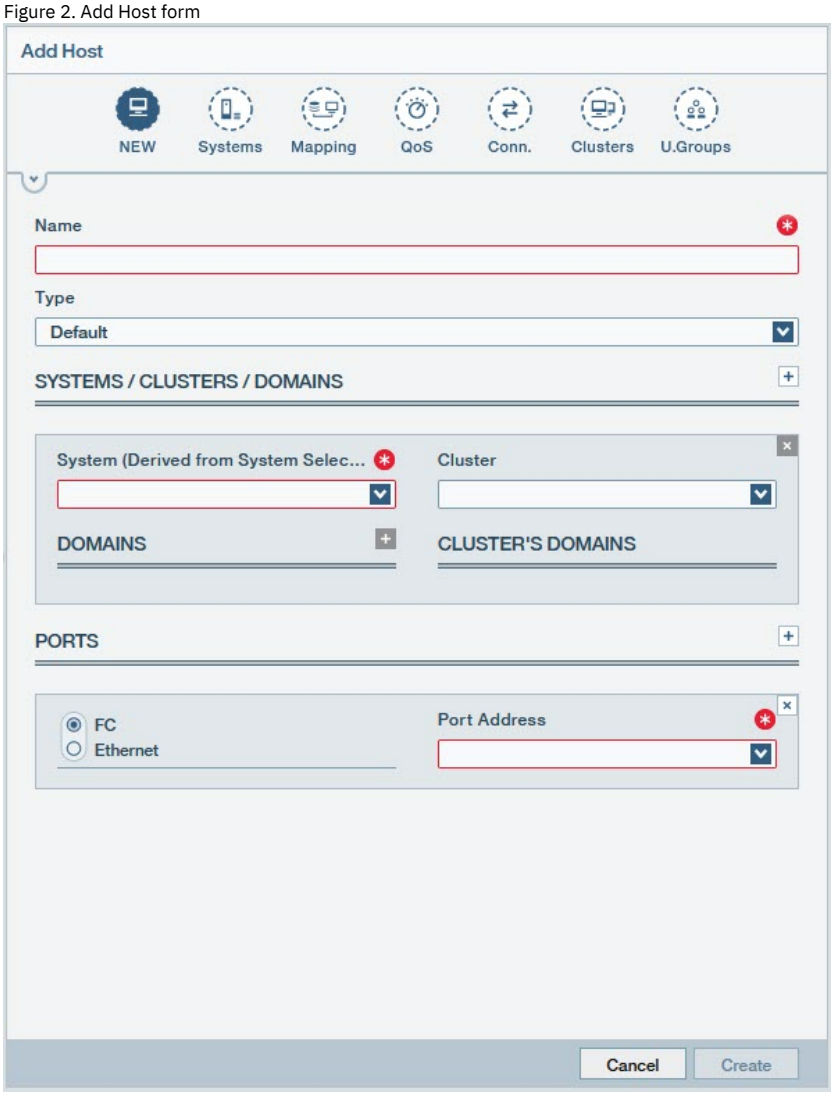

3. Enter the following host details:
   - Name – Define name of the host.
   - Type – Select the host type from the drop down list.
   - PORTS – Select FC or iSCSI to define the port type. Depending on the selected port type, set the port address as follows:
     - For an FC port – Select the World Wide Port Name (WWPN) of the port from the drop down list.
     - For an iSCSI port – Enter the iSCSI Qualified Name (IQN) of the port.
   - System – Select the storage system to which you want to attach the host.
   - Cluster – If the host is a member of a cluster, select the cluster name from the drop down list. Otherwise, leave this option blank.
   - Domain – Select the storage system domain or domains in which this host should be defined.
   - Cluster's Domain – Select the cluster domain or domains in which this host should be defined.
4. Click Apply. The host definition is added to the storage system.

Defining a host from the CLI

1. List a specific host or all hosts:

   ```
   host_list [host=HostName]
   ```

   For more information, see Listing hosts.
2. Define new host and connect it to the storage system:

   ```
   host_define host=HostName [cluster=ClusterName]
   ```

   For more information, see Defining a new host.
3. Add ports to the host, as required:

   ```
   host_add_port host=HostName < fcaddress=wwpn | iscsi_name=iSCSIName [ num_of_visible_targets=num |
   num_of_visible_targets_per_network=num ] >
   ```

   For more information, see Adding a port to a host

Monitoring host connectivity

Monitor the connectivity of attached hosts in real-time through either the Hyper-Scale Manager or the command-line interface (CLI).

Ensure that your host is properly connected to the storage system, by monitoring the host attachment connectivity.

- Monitor using the Hyper-Scale Manager, select the storage system to which the host is connected and navigate to HOSTS & CLUSTERS VIEWS > Host Connectivity.
- Monitor using the XCLI Utility using the **host_connectivity_list** command (see Listing connectivity to hosts).

Configuring Linux host attachment

Configure a host server using a Linux® operating system for host attachment to your storage system.

- **Red Hat Enterprise Linux (RHEL) 8.4 Fibre Channel host attachment**
  Connect your storage system to a Red Hat® Enterprise Linux (RHEL) 8.4 Fibre Channel host operating system.
- **Red Hat Enterprise Linux (RHEL) 8.4 iSCSI host attachment**
  Connect your storage system to a Red Hat Enterprise Linux (RHEL) 8.4 iSCSI host operating system.

Red Hat Enterprise Linux (RHEL) 8.4 Fibre Channel host attachment

Connect your storage system to a Red Hat® Enterprise Linux® (RHEL) 8.4 Fibre Channel host operating system.

Before you begin

From the host, get the host's HBA WWPNs by logging in to the host and running the following commands:

```
[root]# cd /sys/class/fc_host
[root fc_host]# ls -la
total 0
drwxr-xr-x.  2 root root 0 Jul  4 16:01 .
drwxr-xr-x. 68 root root 0 Jul  4 09:38 ..
lrwxrwxrwx.  1 root root 0 Jul  4 16:01 host7 -> ../../devices/pci0000:00/0000:00:02.0/0000:11:00.0/host7/fc_host/host7
lrwxrwxrwx.  1 root root 0 Jul  4 16:01 host8 -> ../../devices/pci0000:00/0000:00:02.0/0000:11:00.1/host8/fc_host/host8
[root fc_host]# cat host7/port_name
0x21000024ff1826d0
[root fc_host]# cat host8/port_name
0x21000024ff1826d1
```

Before configuring on the host side, add the host to the storage system and allocate storage pools and volumes, as needed. This can be done through the Hyper-Scale Manager. For more information, see IBM® Hyper-Scale Manager documentation and the IBM Hyper-Scale Manager for IBM Spectrum® Accelerate Family

Redbooks® publication.
Important: It is necessary to have a zoning between the host and the FlashSystem A9000R system, so the WWPNs can be seen from the storage.

Procedure

Complete the host side configuration by using the following steps.

1. Configure the multipath settings.
   Use the **systemctl status multipathd.service** command and make sure that the Status is *"up"*.
   Important: RHEL has its own native multipath driver, which maps the discovered drives and their paths to the mpath device files in /dev/mapper. Be sure to configure this driver correctly. For more information, see the Configuring device mapper multipath documentation on the Red Hat Customer Portal.
2. Discover the mapped volumes.
   Check that the volumes previously created are detected by the host.
   a. Use the **rescan-scsi-bus.sh** command to scan the new volumes.
      The output shows if the new volumes were discovered.
      For example:

```
[root]# rescan-scsi-bus.sh
Scanning SCSI subsystem for new devices
Scanning host 0 for  SCSI target IDs  0 1 2 3 4 5 6 7, all LUNs
 Scanning for device 0 2 0 0 ...
OLD: Host: scsi0 Channel: 02 Id: 00 Lun: 00
   Vendor: IBM    Model: ServeRAID M5210  Rev: 4.68
   Type:  Direct-Access           ANSI SCSI revision: 05
.....................................................
2 new or changed device(s) found.
0 remapped or resized device(s) found.
0 device(s) removed.
```

   b. Use the **multipath -ll** to show the current multipath configuration.

Red Hat Enterprise Linux (RHEL) 8.4 iSCSI host attachment

Connect your storage system to a Red Hat® Enterprise Linux® (RHEL) 8.4 iSCSI host operating system.

Before you begin

From the host side, install the iSCSI initiator and get the InitiatorName for performing the host connectivity to the storage system.
Note: The *InitiatorName* is necessary to define host connectivity on Hyper-Scale Manager.

1. Install the iscsi-initiator-utils.

```
# sudo  dnf -y install iscsi-initiator-utils
```

2. Get the initiator name.

```
# cat /etc/iscsi/initiatorname.iscsi
InitiatorName=iqn.1994-05.com.redhat:f808c6bed2a
```

Before configuring on the host side, configuring the following on your FlashSystem A9000R system.

These steps can be done through the Hyper-Scale Manager. For more information, see IBM® Hyper-Scale Manager documentation and the IBM Hyper-Scale Manager for IBM Spectrum® Accelerate Family Redbooks® publication.
Important: Connect the host and the storage system so that the IQN can be seen from the storage system.

1. Configure the iSCSI ports by using the Hyper-Scale Manager. For more information, see Configuring iSCSI ports.
2. Add the host to the storage system.
   Note: When creating a new host for an iSCSI host connection, be sure to add the Ethernet port by using the *InitiatorName* in the Port Address parameter.

Procedure

Complete the host side configuration by using the following steps.

1. Identify the network ports that will be used for iSCSI connectivity by using the **ifconfig** command.
2. Define private IP addresses to the network ports identified.
   For example:

```
# ifconfig ens1f0 11.11.110.29 netmask 255.255.254.0
# ifconfig ens1f1 11.11.114.25 netmask 255.255.254.0
```

3. Check whether the changes were applied by using the **ifconfig** command.
4. Discover the iSCSI target from the host.
   Copy or save the IP and IQN values, which will be used in step 5.

```
# iscsiadm -m discovery -t sendtargets -p <ip_address>
```

5. Log in to the discovered targets to establish the sessions from the host. Use the IP and IQN values that were copied in the previous step.

```
# iscsiadm -m node -T <iqn_value> - p <ip_address> -l
```

6. Check the active iSCSI sessions, by running the following command:

```
iscsiadm -m session
```

7. Configure the multipath settings.
   Use the **systemctl status multipathd.service** command and make sure that the Status is *"up"*.
   Important: RHEL has its own native multipath driver, which maps the discovered drives and their paths to the mpath device files in /dev/mapper. Be sure to configure this driver correctly. For more information, see the Configuring device mapper multipath documentation on the Red Hat Customer Portal.
8. Discover the mapped volumes.
   Check that the volumes previously created are detected by the host.
   a. Use the **rescan-scsi-bus.sh** command to scan the new volumes.
      The output shows if the new volumes were discovered.
      For example:

```
[root]# rescan-scsi-bus.sh
Scanning SCSI subsystem for new devices
Scanning host 0 for  SCSI target IDs  0 1 2 3 4 5 6 7, all LUNs
 Scanning for device 0 2 0 0 ...
OLD: Host: scsi0 Channel: 02 Id: 00 Lun: 00
   Vendor: IBM    Model: ServeRAID M5210  Rev: 4.68
   Type:  Direct-Access          ANSI SCSI revision: 05
.....................................................
2 new or changed device(s) found.
0 remapped or resized device(s) found.
0 device(s) removed.
```

   b. Use the **multipath -ll** to show the current multipath configuration.

Configuring Windows host attachment

Configure a host server using a Windows operating system for host attachment to your storage system.

- **Windows Fibre Channel host attachment**
  Connect your storage system to a Windows 2019 or 2022 Fibre Channel host operating system.
- **Windows iSCSI host attachment**
  Connect your storage system to a Windows 2019 or 2022 iSCSI host operating system.

Windows Fibre Channel host attachment

Connect your storage system to a Windows 2019 or 2022 Fibre Channel host operating system.

Before you begin

From the host, get the host's HBA WWPNs by logging in to the host and running the **Get-InitiatorPort** command.

```
PS C:\Users\Administrator> Get-InitiatorPort

InstanceName                                                    NodeAddress      PortAddress      ConnectionType
------------                                                    -----------      -----------      --------------
PCI\VEN_1077&DEV_2261&SUBSYS_02B31077&REV_01\4&3569c1d3&0&0000_0 200034800d700c8a 210034800d700c8a Fibre Channel
PCI\VEN_1077&DEV_2261&SUBSYS_02B31077&REV_01\4&3569c1d3&0&0100_0 200034800d700c8b 210034800d700c8b Fibre Channel
```

Before configuring on the host side, add the host to the storage system and allocate storage pools and volumes, as needed. This can be done through the Hyper-Scale Manager. For more information, see IBM® Hyper-Scale Manager documentation and the IBM Hyper-Scale Manager for IBM Spectrum® Accelerate Family Redbooks® publication.
Important: It is necessary to have a zoning between the host and the FlashSystem A9000R system, so the WWPNs can be seen from the storage.

Procedure

Complete the host side configuration by using the following steps.

1. Enable and configure the multipath settings by using the Configure Multipath I/O for your StorSimple Device article on Microsoft Documentation.
   a. Follow Step 1: Install MPIO on the Windows Server host.
      You can check that the multipath configuration is configured properly by running

```
mpclaim.exe -v c:\multipathconfig.txt
```

      and seeing the current configuration from the generated file.
   b. Follow Step 2: Configure MPIO for StorSimple volumes.
      - If the volume is already mapped to the host, adding the Device Hardware Id can be automatically, as indicated in the article.
      - If needed, the Device Hardware Id can be manually added by clicking Add on the MPIO Devices tab.
        In order to comply with the ID entry requirements, enter the following: `IBM 2810XIV`
        Important: Be sure to capitalize the words *IBM* and *XIV*. In addition, you must insert five spaces after `IBM` and another nine spaces following `2810XIV`.
2. Discover the mapped volumes.
   Check that the volumes previously created are detected by the host.

a. Use the DiskPart **rescan** command to scan the new volumes.
b. Use the **mpclaim.exe -s -d** command to show the current multipath configuration.

Windows iSCSI host attachment

Connect your storage system to a Windows 2019 or 2022 iSCSI host operating system.

Before you begin

From the host side, get the initiator name for performing the host connectivity to the storage system.
Note: The initiator name is necessary to define host connectivity on Hyper-Scale Manager.
Start the Microsoft iSCSI service to get the initiator name from the iSCSI Initiator Properties window.

1. From the host side, use the **iscsipl** command.

   The Microsoft iSCSI dialog opens asking to start the service.

2. Click Yes.
   The iSCSI Initiator Properties window opens.

3. Get the initiator name from the Configuration tab.

Before configuring on the host side, configure the following on your FlashSystem A9000R system.

These steps can be done through the Hyper-Scale Manager. For more information, see IBM® Hyper-Scale Manager documentation and the IBM Hyper-Scale Manager for IBM Spectrum® Accelerate Family Redbooks® publication.
Important: Connect the host and the storage system so that the IQN can be seen from the storage system.

1. Configure the iSCSI ports, using the Hyper-Scale Manager. For more information, see Configuring iSCSI ports.
2. Add the host to the storage system.
   Note: When creating a new host for an iSCSI host connection, be sure to add the Ethernet port by using the *InitiatorName* in the Port Address parameter.

Procedure

Complete the host side configuration by using the following steps.

1. Identify the network ports that will be used for iSCSI connectivity.
   From the Windows Server Datacenter, go to Network settings.
2. From the Network settings page, select the network adapter that needs to be configured.
3. From the Network adapter settings page of the selected adapter, define a private IP address on the selected network port.
   This IP address is used in the next step.
   Check whether the changes were applied from the Network settings page.

4. Discover the iSCSI target from the host.
   From the Targets tab on the iSCSI Initiator Properties window, enter the private IP address, which was defined in the previous step, into the Target field.
5. Click Quick Connect.
6. From the Connect To Target form, add the target to the list of favorites.
   Select Add this connection to the list of Favorite Targets.
7. Click OK.
8. From the Discovery tab, review all active iSCSI sessions.
9. Enable and configure the multipath settings by using the Configure Multipath I/O for your StorSimple Device article on Microsoft Documentation.
   a. Follow Step 1: Install MPIO on the Windows Server host.
      You can check that the multipath configuration is configured properly by running

      ```
      mpclaim.exe -v c:\multipathconfig.txt
      ```

      and seeing the current configuration from the generated file.
   b. Follow Step 2: Configure MPIO for StorSimple volumes.
      - If the volume is already mapped to the host, adding the Device Hardware Id can be automatically, as indicated in the article.
      - If needed, the Device Hardware Id can be manually added by clicking Add on the MPIO Devices tab.
        In order to comply with the ID entry requirements, enter the following: `IBM 2810XIV`
        Important: Be sure to capitalize the words *IBM* and *XIV*. In addition, you must insert five spaces after `IBM` and another nine spaces following `2810XIV`.
10. Discover the mapped volumes.
    Check that the volumes previously created are detected by the host.
    a. Use the DiskPart **rescan** command to scan the new volumes.
    b. Use the **mpclaim.exe -s -d** command to show the current multipath configuration.

Configuring Solaris host attachment

Connect your storage system to a Solaris 10 or 11 Fibre Channel (FC) host operating system.

Before you begin

From the host, get the host's HBA WWPNs by logging in to the host and running the **fcinfo hba-port** command.
Before configuring on the host side, add the host to the storage system and allocate storage pools and volumes, as needed. This can be done through the Hyper-Scale Manager. For more information, see IBM® Hyper-Scale Manager documentation and the IBM Hyper-Scale Manager for IBM Spectrum® Accelerate Family Redbooks® publication.
Important: It is necessary to have a zoning between the host and the FlashSystem A9000R system, so the WWPNs can be seen from the storage.

Procedure

Complete the host side configuration by using the following steps.

1. Solaris multipathing is done locally, through Solaris I/O multipathing.
   Enable multipathing by using the one of the following instructions, within the relevant Solaris documentation:
   - **Oracle Solaris 10 Documentation Library:** SAN Configuration and Multipathing
   - **Oracle Solaris 11 Documentation Library:** SAN Configuration and Multipathing
2. Check that multipathing is activated.
   Check the mpxio-disable parameter in the configuration file. If `mpxio-disable="yes"`, the driver is not activated and multipathing must be enabled, as described in step 3. The driver is not activated by default.

   Example for Solaris 10

   ```
   egrep "^mpxio-disable=" /kernel/drv/*.conf
   /etc/driver/drv/fp.conf:mpxio-disable="no";
   /etc/driver/drv/lsc.conf:mpxio-disable="no";
   /etc/driver/drv/mpt_sas.conf:mpxio-disable="no";
   /etc/driver/drv/mpt.conf:mpxio-disable="no";
   /etc/driver/drv/pmcs.conf:mpxio-disable="no";
   ```

   Example for Solaris 11

   ```
   egrep "^mpxio-disable=" /etc/driver/drv/*.conf
   /etc/driver/drv/fp.conf:mpxio-disable="no";
   /etc/driver/drv/lsc.conf:mpxio-disable="no";
   /etc/driver/drv/mpt_sas.conf:mpxio-disable="no";
   /etc/driver/drv/mpt.conf:mpxio-disable="no";
   /etc/driver/drv/pmcs.conf:mpxio-disable="no";
   ```

3. Optional: Enable multipathing.
   Multipathing needs to be enabled if `mpxio-disable="yes"` and in some cases after inserting a new HBA.
   Use the **stmsboot -e** command on each HBA controller on the host.
   Important: Rebooting the system is required in order for the changes to take place.

   ```
   stmsboot -e
   WARNING: stmsboot operates on each supported multipath-capable controller
                detected in a host. In your system, these controllers are

      /devices/pci@780/pci@0/pci@8/SUNW,qlc@0/fp@0,0
      /devices/pci@780/pci@0/pci@8/SUNW,qlc@0,1/fp@0,0
      /devices/pci@7c0/pci@0/pci@1/pci@0,2/LSILogic,sas@1
      /devices/pci@7c0/pci@0/pci@1/pci@0,2/LSILogic,sas@1
      /devices/pci@7c0/pci@0/pci@1/pci@0,2/LSILogic,sas@2
      /devices/pci@7c0/pci@0/pci@9/LSILogic,sas@0
      /devices/pci@7c0/pci@0/pci@9/LSILogic,sas@0

      If you do NOT wish to operate on these controllers, please quit stmsboot
      and re-invoke with -D { fp | mpt } to specify which controllers you wish
      to modify your multipathing configuration for.

      Do you wish to continue? [y/n] (default: y) y
      Checking mpxio status for driver fp
      Checking mpxio status for driver mpt
      WARNING: This operation will require a reboot.
      Do you want to continue ? [y/n] (default: y) y
      The changes will come into effect after rebooting the system.
      Reboot the system now ? [y/n] (default: y) y
   ```

4. Optional: Disabling multipathing for Veritas configuration.
   Note: Veritas configuration requires native MPxIO to be disabled, as Veritas DMP (Dynamic Multipathing) with MPxIO configuration is not supported.
   Disable multipathing, by using the **stmsboot -d** command. This runs on each HBA controller on the host.
   Important: Rebooting the system is required in order for the changes to take place.

   ```
   stmsboot -d
   WARNING: stmsboot operates on each supported multipath-capable controller
                detected in a host. In your system, these controllers are

      /devices/pci@780/pci@0/pci@8/SUNW,qlc@0/fp@0,0
      /devices/pci@780/pci@0/pci@8/SUNW,qlc@0,1/fp@0,0
      /devices/pci@7c0/pci@0/pci@1/pci@0,2/LSILogic,sas@1
      /devices/pci@7c0/pci@0/pci@1/pci@0,2/LSILogic,sas@1
      /devices/pci@7c0/pci@0/pci@1/pci@0,2/LSILogic,sas@2
      /devices/pci@7c0/pci@0/pci@9/LSILogic,sas@0
      /devices/pci@7c0/pci@0/pci@9/LSILogic,sas@0

      If you do NOT wish to operate on these controllers, please quit stmsboot
      and re-invoke with -D { fp | mpt } to specify which controllers you wish
   ```

```
    to modify your multipathing configuration for.

    Do you wish to continue? [y/n] (default: y) y
    Checking mpxio status for driver fp
    Checking mpxio status for driver mpt
    WARNING: This operation will require a reboot.
    Do you want to continue ? [y/n] (default: y) y
    The changes will come into effect after rebooting the system.
    Reboot the system now ? [y/n] (default: y) y
```

5. Discover the mapped volumes.
   Run Solaris OS commands, or external vendor commands, to detect the newly attached LUNs. The following is an example of a general procedure to detect new FC devices.

```
cfgadm -al -o show_FCP_dev
Ap_Id                          Type         Receptacle   Occupant     Condition
c2                             fc-fabric    connected    configured   unknown
c2::10000000c9cc0489           unknown      connected    unconfigured unknown
c2::500173800ce80111,0         array-ctrl   connected    configured   unknown
c2::500173800ce80111,1         disk         connected    configured   unknown
c2::500173800ce80111,2         disk         connected    configured   unknown
c2::500173800ce80121,0         array-ctrl   connected    configured   unknown
c2::500173800ce80121,1         disk         connected    configured   unknown
c2::500173800ce80121,2         disk         connected    configured   unknown
c2::500173800cea0111,0         array-ctrl   connected    configured   unknown
c2::500173800cea0111,1         disk         connected    configured   unknown
c2::500173800cea0111,2         disk         connected    configured   unknown
c2::500173800cea0123,0         array-ctrl   connected    configured   unknown
c2::500173800cea0123,1         disk         connected    configured   unknown
c2::500173800cea0123,2         disk         connected    configured   unknown
```

6. Format the volumes and label each of the devices that need to be configured.
   For example:
   Note: The exact output of the **format** command might vary is Solaris MPxIO is enabled.

```
format
Searching for disks...done


AVAILABLE DISK SELECTIONS:
       0. c5t5000CCA03C43B8A9d0 <HITACHI-H106030SDSUN300G-A2B0-279.40GB>
          /pci@400/pci@1/pci@0/pci@4/scsi@0/iport@1/disk@w5000cca03c43b8a9,0
          /dev/chassis/SYS/HDD0/disk
       1. c7t5000CCA03C42E74Dd0 <HITACHI-H106030SDSUN300G-A2B0 cyl 46873 alt 2 hd 20 sec 625>
          /pci@400/pci@2/pci@0/pci@4/scsi@0/iport@1/disk@w5000cca03c42e74d,0
          /dev/chassis/SYS/HDD4/disk
       2. c2t500173800CE80111d1 <IBM-2810XIV-0000 cyl 117 alt 2 hd 128 sec 128>
          /pci@400/pci@2/pci@0/pci@c/SUNW,qlc@0/fp@0,0/ssd@w500173800ce80111,1
       3. c2t500173800CE80111d2 <IBM-2810XIV-0000 cyl 117 alt 2 hd 128 sec 128>
          /pci@400/pci@2/pci@0/pci@c/SUNW,qlc@0/fp@0,0/ssd@w500173800ce80111,2
       4. c2t500173800CE80111d3 <IBM-2810XIV-0000 cyl 117 alt 2 hd 128 sec 128>
          /pci@400/pci@2/pci@0/pci@c/SUNW,qlc@0/fp@0,0/ssd@w500173800ce80111,3
Specify disk (enter its number): 2
selecting c2t500173800CE80111d1
Disk not labeled.  Label it now?  yes
format> verify
format> q
```

7. List all multipath devices on Solaris by using the appropriate commands for your Solaris system.

   Listing multipath devices for Solaris MPxIO systems
   Use the **mpathadm** utility and run the **mpathadm list lu** command.
   Listing multipath devices for Veritas DMP systems
   a. Either rescan or devices on the operating system level by using the **cfgadm** command and then run the **vxdctl -f enable** command
      or
      use the **dmpdr** utility, run the following command:

```
/usr/lib/vxvm/voladm.d/bin/dmpdr -o refresh
```

   b. List all multipath devices on Veritas, using the **vxdisk list** utility.

Configuring key servers

Encryption key servers create and manage encryption keys that are used by the system. In environments with many systems, key servers distribute keys remotely without requiring physical access to the systems.

Thales CipherTrust Manager and SafeNet KeySecure key servers support KMIP and create keys on demand, and then shares them with the other clustered key servers, providing redundant access. The system supports different types of configurations on key servers. The following configurations are supported:

- Thales CipherTrust Manager and KeySecure key servers use an active-active model, where multiple key servers are used to provide redundancy. In these configurations, one key server must be specified as the primary key server. The primary key server is the key server that the system uses when you create any new encryption keys. The key is immediately replicated to the other key servers in the cluster. All of the key servers that are defined on the system can be used to retrieve keys. Although it is possible to configure a single key server instance, two key servers are recommended to ensure availability of keys, if one key server experiences an outage.

- The system supports up to four key servers. If the system is accessing multiple key servers, they need to belong to the same cluster of key servers.

For the latest supported Thales CipherTrust Manager and Gemalto SafeNet KeySecure key servers, see Encryption key management schemes.
Note: Be sure to read the Release Notes for your Thales CipherTrust Manager product version, on the Thales CPL Documentation Hub. For best security practices and to help prevent potential downtime and configuration issues, be sure to use only supported TLS versions and read the recommendations for secure initialization vector (IV) for AES.
For more information, see the following resources:

- Data-at-rest encryption
- Data-at-rest Encryption for the IBM Spectrum Accelerate Family Redpaper publication (REDP-5402)

- **Thales CipherTrust Manager k170v v2.3.x and 2.6.x migration to client certificate authentication key**
  Migrate the global key management scheme to client certificate authentication key with Thales CipherTrust Manager server k170v 2.3.x or 2.6.x on your storage system.
- **Thales CipherTrust Manager k170v v2.2.x migration to client certificate authentication key**
  Migrate the global key management scheme to client certificate authentication key with Thales CipherTrust Manager server k170v 2.2.x on your storage system.
- **Thales CipherTrust Manager k170v v2.0.0 migration to client certificate authentication key**
  Migrate the global key management scheme to client certificate authentication key with Thales CipherTrust Manager server k170v 2.0.0 on your storage system.
- **Thales SafeNet KeySecure 8.11.1 migration to client certificate authentication key**
  Migrate the global key management scheme to client certificate authentication key with Thales SafeNet KeySecure 8.11.1 on your storage system.
- **Thales SafeNet KeySecure 8.11.1 client certificate authentication for key management**
  Configure an external encryption key management scheme with Thales SafeNet KeySecure 8.11.1 Client Certificate Authentication on your storage system.

Thales CipherTrust Manager k170v v2.3.x and 2.6.x migration to client certificate authentication key

Migrate the global key management scheme to client certificate authentication key with Thales CipherTrust Manager server k170v 2.3.x or 2.6.x on your storage system.

The information in this section:

- Does not include installation of the key server.
- Does not cover all available configuration options. It contains the basic configuration to allow encryption between a FlashSystem A9000R storage system and the CipherTrust Manager key server.
- Uses the storage system management port example IP address `9.11.80.61` and the CipherTrust Manager server address `9.11.110.162`.
  Replace these addresses with the correct IP addresses of your FlashSystem A9000R storage system and of your key server.

- Uses the FlashSystem A9000R machine serial number as a prefix when naming PEM files. For example, if the file is saved as 1340006-ca.pem, *134006* indicates the machine serial number of the storage array. To get the serial number of a FlashSystem A9000R storage array, use the **config_get** CLI command (see Displaying the values of configuration parameters).

Known issues

When using Thales CipherTrust Manager with FlashSystem A9000R, and disabling encryption, it may fail with `ENCRYPT_UNABLE_TO_UPDATE_KEY_DURING_DEACTIVATE_ON_KEYSERVER`. This happens if the system was already encrypt enabled, then disabled and enabled again. This can also happen if the key was in Pre-Active state on the key server when encrypt enabling occurred and if this state was not changed later to 'Active'. In order to prevent this issue from occurring, before encryption, enable the key on the key server side. The key on the key server cannot be deactivated or re-activated if it was already deactivated once. If it is deactivated, a new key will be needed on the key server side.

- **Enabling encryption on the storage system**
  Enable encryption on the storage system side with the CipherTrust Manager global key authentication.
- **Migrating from global key to client certificate authentication**
  Migrate from global key to Client Certificate Authentication.
- **CipherTrust Manager cluster configuration**
  Provide additional protection and maintain the key server cluster by creating a backup key server.
- **Using an external CA certificate**
  Configure an external CA on the CipherTrust Manager and on the storage system.

Enabling encryption on the storage system

Enable encryption on the storage system side with the CipherTrust Manager global key authentication.

Preconfigure the storage system by using the following steps:

1. Verify that the storage system is encryption-ready by running the **state_list** command (see Listing the operational state).
   If encryption is not supported, contact IBM Support.

2. Check if any certificates are defined on the system, using the **pki_list** command (see Listing PKI items).
   Note: Having certificates that are predefined on the system does not cause any issues with the configuration.

Configuring the key server

All procedures that are detailed here take place on the Thales CipherTrust Manager, unless otherwise stated.

1. Create a local Certificate Authority (CA).
   Note:
   - This step might not be necessary if a CA is already in the Local Certificate Authority List If using an External CA, see Using an external CA certificate.
   - This PEM text example and its decoded X.509 are presented for illustrative purposes only and do not completely fit the values shown in the associated screen captures. For example, there may be inconsistencies in the Common Name (CN) and the Certificate Name field values.

   a. Navigate to CA.
   b. From the Local Certificate Authorities section, select New Local CA.
      Important: Use a strong enough key length and algorithm, such as RSA-2048, according to your company's security policies.
   c. Fill in the relevant information and click Create Local CA. The certificate will appear under Pending CAs in the Local Certificate Authority List.
   d. Select the Pending CA and click Self Sign. The CA now appears in the Local Certificate Authority List.
   e. Click on the ellipses icon (...), next to the CA certificate, and select Download.
      This file will be used in later steps to create the client chained certificate. (See step 5 of Configuring the storage system.)

   f. Save the root (local) CA certificate in PEM format. For example, 1340006-ca.pem.
      The PEM format encodes an X.509 certificate.

   g. Check that the root certificate can sign other certificates. Use the following command:

      ```
      # openssl x509 -in 1340006-ca.pem -inform pem -text -noout
      ```

      Note: The following are indications of a root certificate which can sign other certificates:
      - The Issuer and the Subject are the same.
      - There is an X.509 extension indicating it is a CA (`CA:TRUE`).

2. Configure KMIP.
   a. Navigate to Access Management > Registration Tokens.
   b. Click + New Registration Token.
      i. Assign a name for the new token in the Name Prefix field and click Next.
      ii. Select the Local CA that will be used and click Create Token.
      iii. Copy the token and click Done.
         Note: This token will be used in the next step.
   c. Navigate to Admin Setting > Interfaces.
      i. If the KMIP interface does not exist, it needs to be created.
      ii. Click on the ellipses icon (...) and select Edit.
      iii. Configure the following if not already set:
         - Select Enable hard delete and Auto Registration.
         - In the Registration Token field, paste the token that was copied in step 2.b.iii.
         - From the Mode drop-down list, select TLS, verify client cert, allow anonymous logins.
         - From the Local CA for Automatic Server Certificate Generation drop-down list, select your local CA.
         - Add in all local trusted CAs from the Local Trusted CAs drop-down list and click the add (+) button.

      iv. Click Update.
   d. Restart the NAE-KMIP service.
      i. Navigate to Admin Settings > System Services.
      ii. Restart the NAE-KMIP service by clicking Restart.
         Important: This is required to set the CA for KMIP configuration.
3. Retrieve the KMIP server certificate.
   a. Navigate to CA.
      i. Select the newly created Local CA.
         The new kmip.keysecure.local certificate is within the CA.

      ii. Click on the ellipses icon (...) of the kmip.keysecure.local certificate and click Download.
      iii. Save the server certificate in PEM format. For example, key_server_9.11.110.162.pem.
         Note: Up to four key servers can be configured at the same time on each FlashSystem A9000R storage array and each one must be configured with its own certificate that is signed by the same CA certificate. Therefore, be sure to identify each server PEM file with the IP address or FQDN of the specific key server for easy identification. When using the **encrypt_keyserver_define** command, be sure to use all existing signed storage array certificates.
         The PEM file is used in 6 of Configuring the storage system.

         The PEM format encodes an X.509 certificate.

      iv. Check that the root certificate can sign other certificates.
         Use the following command:

         ```
         # openssl x509 -in key_server_9.11.110.162.pem -inform pem -text -noout
         ```

         Note: The following are indications of a signed certificate which cannot sign other certificates:
         - The Issuer and the Subject are *not* the same.
         - There is an X.509 extension indicating it is *not* a CA (`CA:FALSE`).

4. Create an AES encryption key.
   a. Navigate to Keys.
   b. To add an AES encryption key, click Create a New Key.
   c. Set the Create Key fields.
      - Fill in the Key Name. The key name is the machine model and serial number of your storage system. For example, `425-1340006`, where *425* is the machine model and *1340006* is the serial number.
      - From the Algorithm drop-down list, select AES.
      - From the Size drop-down list, select 256.

Click Create.

5. Add the Key Users group.
   a. Set Added Only to off.
   b. In the search groups field, find Key Users.
6. Set Exportable and Deletable to on.
7. On the Default tab, within the KEY ACCESS form, set the Key Owner to global from the drop-down list.
8. On the KMIP tab, navigate to the KMIP | ALIASES & ALTERNATIVE NAMES form.
   a. In the Aliases field, write in the Key Name and click the plus icon (+) to add it.
   b. Click Update.

Configuring the storage system

1. Creating the Certificate Signing Request (CSR) for the KMIP client certificate.
   Create the request on the storage system.
   a. Create the Certificate Signing Request (CSR) on the storage system side.
      This generates a certificate, which is later used to sign with local CA in CipherTrust Manager Server (see step 2 of Configuring the storage system).

      Define the following parameters in the **pki_generate_private_key_and_csr** command (see Generating a private key and CSR):
      - Name is the certificate's symbolic name.
      - Subject is the subject name for the generated certificate request.

      The argument must be formatted as **/type0=value0/type1=value1/type2=...** .
      Include the email of the person who needs to be notified of any certificate or security issues.

      For example,

      ```
      pki_generate_private_key_and_csr name=CipherTrust subject="/C=US/ST=Arizona/L=Tucson/O=IBM/OU=IBM-
      lab/CN=secadmin1/emailAddress=name@company.com"
      CSR
      ```

      where *CipherTrust* is the name and *secadmin1* is the CN (Common Name) username, with the storage array located in Tuscon, Arizona.
      The storage system certificate is displayed.

   b. Copy the entire text of the storage system SSL Certificate Signing Request (CSR), including the **`-----BEGIN CERTIFICATE REQUEST-----`** and **`-----END CERTIFICATE REQUEST-----`**.
   c. Verify the PKI items, using the **pki_list** command (see Listing PKI items).

      ```
      pki_list
      Name          Fingerprint                          Has signed certificate   Services
      ----------    ---------------------------------- ------------------------ ----------------------------
      1340006       34b52765fbca8c1972931d63e581050c    yes                      XCLI,CIM,IPSEC,KMIP,QUORUM
      CipherTrust    4fd2da6d2ac67e966ae02ad59f0c9147   no                       NONE
      ```

2. Sign the storage system Certificate Signing Request (CSR).
   Use a CipherTrust Manager key secure local CA to sign the Client Certificate Request.
   a. Navigate to CA.
   b. Select a local CA created with the procedures in Creation of a Local CA section.
   c. Select Upload and Sign CSR.
   d. From the Upload Externally Generated CSR form:
      i. Paste the certificate request into the Certificate Request field.
      ii. Select the Client as the Certificate Purpose.
      iii. Specify a Certificate Duration.
   e. Once the form is complete, click Issue Certificate.
3. Download the signed storage system client certificate PEM file.
   a. Click on the ellipses icon (...) next to the newly signed certificate and select Download.
   b. Save the signed client certificate in PEM format. For example: 1340006-client.pem
      The PEM format encodes an X.509 certificate.

4. Check that the client certificate cannot sign other certificates (that it is a leaf certificate).
   Use the following command:

   ```
   # openssl x509 -in 1340006-client.pem -inform pem -text -noout
   ```

   Note: The following are indications of a signed certificate which cannot sign other certificates:
   - The Issuer and the Subject are *not* the same.
   - There is an X.509 extension indicating it is *not* a CA (**CA:FALSE**).
5. Create a chained client certificate PEM file.
   These steps must be done on the Linux host where the two files were downloaded from the key server.

   Use the local root CA PEM file, created and downloaded in step 1 of Configuring the key server (in these examples, 1340006-ca.pem) and the signed storage client certificate PEM file, downloaded in step 3 of Configuring the storage system (in these examples, 1340006-client.pem).

   Ensure that the file used in the CLI command exists on the host where the XCLI Utility is installed. If the files downloaded from CipherTrust Manager are on a different Linux host, the files must be copied to the host where the utility is installed. If needed, a utility, such as SCP (secure copy), may be used to copy the file.

   a. Create a chained client certificate PEM file by combining the local and client certificates.
      Before creating the chained certificate, the certificates can be viewed by using the **cat <certificatename>.pem** command.

      Create the chained certificate using the following command:
      Important: The chained certificate starts with the leaf (for example, the storage system client certificate) and ends with the root certificate (the local CA certificate).

Verify that the chain is setup in the correct order in order to avoid configuration issues.

For example:

```
cat 1340006-client.pem 1340006-ca.pem >1340006-chained-client.pem
```

b. Install the chained client certificate on the storage system and update the certificate's related services.
From the storage system XCLI Utility:
   i. Import the signed PEM certificate, using the **pki_set_pem** command (see Importing a signed certificate).

```
pki_set_pem services=none certificate="$(tr "\n" "*" </root/IBM_Storage_Extended_CLI/1340006-chained-client)"
```

   ii. List the PKI items, using the **pki_list** command.
   iii. Remove the KMIP service from the services list of the chained client certificate, using the following `pki_update` CLI command (see Updating a PKI certificate or services).
   The expected output is `Command completed successfully.`

   iv. Add the KMIP service to the CipherTrust certificate, using the **pki_update** command.

```
pki_update name=csrforks services=KMIP
```

   v. Verify that all PKI items are properly updated, using the **pki_list** command.

6. Create a chained server certificate PEM file.
These steps must be done on the Linux host where the two files were downloaded from the key server.

   Use the local root CA PEM file, created and downloaded in steps 1 of Configuring the key server and the KMIP server certificate (in these examples, 1340006-ca.pem), retrieved in step 3 of Configuring the key server (in these examples, key_server_9.11.110.162.pem).

   Ensure that the file used in the CLI command exists on the host where the XCLI Utility is installed. If the files downloaded from CipherTrust Manager are on a different Linux host, the files must be copied to the host where the utility is installed. If needed, a utility, such as SCP (secure copy), may be used to copy the file.

   Create a chained server certificate PEM file by combining the local and server certificates. Before creating the chained certificate, the certificates can be viewed by using the **cat <certificatename>.pem** command.

   Create the chained certificate using the following command:
   Important: The chained certificate starts with the leaf (for example, the server certificate) and ends with the root certificate (the local CA certificate).
   Verify that the chain is setup in the correct order in order to avoid configuration issues.

   For example:

```
cat key_server_9.11.110.162.pem 1340006-ca.pem > 1340006-chained-key_server_9.11.110.162.pem
```

7. Configure the CipherTrust Manager key server on the storage system.
On the host that will be used to run CLI commands, make sure that the storage system chained server certificate (for example, 1340006-chained-key_server_9.11.110.162.pem.) created in step 6 of Configuring the storage system exists.

   a. Perform the following CLI commands from the storage system:
   Note:
   - Use a Linux machine to perform this step as the storage system expects the certificate data to be streamed.
   - The default port is 5696. This can be changed if necessary.

   i. **encrypt_keyserver_define** with the following parameters defined (see Defining a key server):
      - name: `globalkey`
      - ipv4: the Thales key server
      - master: `yes`
      - keyserver_type: `KeySecure`
      - certificate: the chained server PEM certificate

      For example:

```
encrypt_keyserver_define name=globalkey ipv4=(THALES KeyServer IP) keyserver_type=KeySecure master=yes
certificate="$(tr "\n" "*" </root/IBM_Storage_Extended_CLI/1340006-chained-key_server_9.11.110.162.pem)"
```

      Note: When using the **encrypt_keyserver_define** command, be sure to use all existing signed storage array certificates.
   ii. Check the key server, using the **encrypt_keyserver_list** command (see Displaying key server status).
   Note: At this point, when encryption on the storage system has still not be enabled, it is okay to see *NOKEY* in the App/Key Status column. This means that there is still no key encrypting the data.

Creating recovery keys

Best practice recommends that recovery keys be configured. When recovery keys are configured, be sure that the command below includes **`recovery_keys=yes`**.

.For more information regarding recovery keys and how to generate them, see the "Recovery key use and maintenance" section in the *Data-at-rest Encryption for the IBM Spectrum Accelerate Family* Redpaper (REDP-5402).

In order to create recovery keys, at least two Security Admin users must be defined on the storage system.

1. Define at least two Security Admin users on the storage system.
   a. Define each of the Security Admin users, using the **user_define** command (see Defining a new user). For example:

```
user_define user=secadmin1 category=securityadmin password=<password> password_verify=<password>
user_define user=secadmin2 category=securityadmin password=<password> password_verify=<password>
```

b. Verify the creation of both Security Admins, using the **user_list** command (seeListing users).

2. Create the recovery keys.
   Create two recovery keys, one per Security Admin user.
   a. Define the Security Admin users who are to be receiving the recovery key, using the **encrypt_recovery_key_generate** command (see Generating recovery keys).
      For example:

      ```
      encrypt_recovery_key_generate users=secadmin1,secadmin2  min_req=2
      ```

   b. Log in with each Security Admin credential and generate the recovery key, using the **encrypt_recovery_key_get** command (see Retrieving the security administrator's recovery key).
      The following is an example of a full recovery key:

      ```
      62807CB1902AM074EDLA4EV8F0C574E40A1564F55570CDEEBED37BC3876789
      ```

3. Each Security Admin must verify their recovery key, using the **encrypt_recovery_key_verify** command (see Recovering key verification).

Activating the encryption

1. Update the key server, using the **encrypt_enable recovery_keys=yes -y** command (see Enabling encryption).
2. Run the **state_list** command to check that encryption is *Enabled* (see Listing the operational state). If the status is *Enabling*, run the command again, until status is *Enabled*.
3. Check the key server, using the **encrypt_keyserver_list** command (see Displaying key server status).
   The App/Key Status column should show an *Active* status.

Verifying the status

After the encryption keys have been successfully enabled, the flash enclosures and vault devices start changing from a *Ready* to an *Enrolled* state.

- Verify the encryption state of the flash enclosures and vault devices, by running the following CLI commands as a Security Administrator:
  Note: There are corresponding events visible when using the **event_list** command (see Listing events).
  - **flash_enclosure_list -S** (see Listing flash enclosure status)
  - **vault_device_list -S** (see Listing vault devices in the system)

Migrating from global key to client certificate authentication

Migrate from global key to Client Certificate Authentication.

All procedures detailed in this section take place within the Thales CipherTrust Manager Management Console, unless otherwise stated.

Creating a Local User on CipherTrust Manager

If a user has already been created to be used in the Client Certificate Authentication of other devices, skip this section, and proceed to Editing the AES encryption key on CipherTrust Manager.

If a unique user for each device is to be used, this step is mandatory for each and for each FlashSystem A9000R storage array.

If the Certificate Signing Request (CSR) is created with multiple Common Names (CNs), create a user for each CN.

1. Define a user, from Access Management > Users.
   a. Click Create New User.
   b. Define a Username. This username is also used as the Storage Security user.
   c. Fill in the form for the Email, Password, and Certificate Subject Distinguished Name fields.
   d. Select Allow user to login using certificate.
   e. Click Create.
2. Repeat step 1 to create a second user.
3. Configure KMIP.

Editing the AES encryption key on CipherTrust Manager

1. From the Keys menu, select the key that was created in 4 of Configuring the key server, with your storage system model and serial number.
2. In the KEY ACCESS form, use the Key Owner drop-down list to select the user created previously in Creating a Local User on CipherTrust Manager.
   Note: When creating a CSR and using multiple CNs, the key owner is last on the CN list.
3. Scroll to the bottom of the page and select Update.
   The new owner is now reflected on the Keys main page.

Editing the KMIP key server configuration

1. Navigate to Admin Settings > Interfaces.
2. Click on the ellipses icon (...) of the kmip interface and click Edit.
3. From the Mode drop-down list, select TLS, verify client cert, user name taken from client cert, auth request is optional.
4. From the Username Location in Certificate drop-down list, select CN.
5. Click Update.

6. Navigate to Admin Settings > Services.
7. Reboot the NAE-KMIP service by clicking Restart.

Verifying and configuring the storage system

This section assumes that encryption is enabled with the global key user. If this is not the case, refer to Configuring the storage system and enable the global key user before proceeding.

1. Check the CipherTrust Manager Key Server configuration on your storage system.
   To check the configuration use the following CLI commands:
   a. **encrypt_keyserver_check_status** (see Checking key server status).
   b. **encrypt_keyserver_list** (see Displaying key server status).
      If the App/Key Status *does not* show the *Active* status, continue to step 2.

      If the status is *Active*, the migration and configuration is complete. For further protection, create a backup key server (see CipherTrust Manager cluster configuration).

2. Change the Key Status to *Active*.
   a. Navigate to KMIP > Registered Clients.
   b. Click on the ellipses icon (...) for each of the registered clients, and select Delete for each instance.
   c. Navigate to Admin Settings > Services and restart the NAE-KMIP service, by clicking Restart.
3. Recheck the Key Status.
   a. Use the commands from step 1.
      The App/Key Status should now be *Active*.

   b. Verify the status again through the storage system events log. For example:

```
2021-03-04 16:09:44 849272 ENCRYPT_UNABLE_TO_RETRIEVE_KEY_FROM_KEYSERVER Failed to retrieve key from key server
'KS' via IPv4 on module 1. Please verify that the key server type and version are supported. If so, please
check its status.
2021-03-04 16:09:44 849273 ENCRYPTION_SKMIP_ERROR           Module 3 reported KS returned error:1 - 1
2021-03-04 16:09:44 849274 ENCRYPT_UNABLE_TO_RETRIEVE_KEY_FROM_KEYSERVER Failed to retrieve key from key server
'KS' via IPv4 on module 3. Please verify that the key server type and version are supported. If so, please
check its status.
2021-03-04 16:09:44 849275 ENCRYPT_UNABLE_TO_RETRIEVE_KEY_FROM_KEYSERVER Failed to retrieve key from key server
'KS' via IPv4 on module 2. Please verify that the key server type and version are supported. If so, please
check its status.
2021-03-04 16:09:44 849276 EMAIL_NOT_SENT                   Sending event
ENCRYPT_UNABLE_TO_RETRIEVE_KEY_FROM_KEYSERVER (849272) to Event Center via us.relay.ibm.com was  waived because
of failed SMTP gateway. It will be not be used until 2021-03-04 16:11:27.
2021-03-04 16:09:44 849277 EMAIL_NOT_SENT                   Sending event ENCRYPTION_SKMIP_ERROR (849273) to
Event Center via us.relay.ibm.com was  waived because of failed SMTP gateway. It will be not be used until
2021-03-04 16:11:27.
2021-03-04 16:09:44 849278 EMAIL_NOT_SENT                   Sending event
ENCRYPT_UNABLE_TO_RETRIEVE_KEY_FROM_KEYSERVER (849274) to Event Center via us.relay.ibm.com was  waived because
of failed SMTP gateway. It will be not be used until 2021-03-04 16:11:27.
2021-03-04 16:09:44 849279 EMAIL_NOT_SENT                   Sending event
ENCRYPT_UNABLE_TO_RETRIEVE_KEY_FROM_KEYSERVER (849275) to Event Center via us.relay.ibm.com was  waived because
of failed SMTP gateway. It will be not be used until 2021-03-04 16:11:27.
2021-03-04 16:10:58 849280 ENCRYPT_KEYSERVER_EDITED         Details of key server '9.11.54.106' were modified.
```

CipherTrust Manager cluster configuration

Provide additional protection and maintain the key server cluster by creating a backup key server.

The backup key server is important in order to avoid a single point of failure (SPOF) in case the key server fails. When a backup key server is in place and a keyserver failure occurs, the backup that is clustered with it can keep running. The preferred practice is that the backup is located in a different physical location than the primary server.

Configuring the cluster

This section assumes that there are two functioning key servers.

Use the CipherTrust Manager Management Console to complete the following procedure.

1. Navigate to Admin Settings > Cluster.
2. Select Manage Cluster > Create Cluster.
3. Configure the local key server and click Save.
4. From the Cluster Configuration page, click Manage Cluster.
5. Click Add Node.
6. Configure the second key server.
   a. From the Public address of the new node field, configure the IP address of the second key server.
   b. Click Add Node. This opens the second CipherTrust Manager key server.
   c. Click Join. This guides you through the process of clustering the two key servers.
   d. Click Finish! to complete the process.
      The result displays both keys servers in the cluster in a *Ready* state.

Configuring two key servers on the storage system

This section assumes that one key server is already configured.

Define the second key server, using the **encrypt_keyserver_define** command (see Defining a key server).

Use the same chained server certificate PEM file used in the first key server configuration and ensure that to set `master=no`.
For example:

```
encrypt_keyserver_define name=a9 ipv4=(THALES KeyServer IP) keyserver_type=KeySecure master=no certificate="$(tr "\n" "*"
</root/IBM_Storage_Extended_CLI/1340006-chained-key_server_9.11.110.162.pem)"
```

Using an external CA certificate

Configure an external CA on the CipherTrust Manager and on the storage system.

Configuring the external CA on CipherTrust

Be sure that the Certificate Signing Request (CSR) with the external CA is signed with the TLS Web Client Authentication option enabled.

Use the CipherTrust Manager Management Console to complete the following procedure.

1. Upload the external certificate.
   a. From the CA form:
      - **For k170v 2.3.x:** Select Upload External Certificate.
      - **For k170v 2.6.x:** Select the External tab and then Add External CA.
   b. Paste the external certificate in PEM format.
   c. Click Upload.
2. Add the CA to the trusted list for the KMIP interface.
   a. Navigate to Admin Settings > Interfaces.
   b. Click on the ellipses icon (...) for the KMIP interface and select Edit.
   c. From the External Trusted CAs drop-down list, select the external CA that was uploaded in step 1.
   d. Click on the plus icon (+) to add the CA. This adds the CA into the trusted list for the KMIP interface.
   e. Click Update.
   f. Navigate to Admin Settings > Services and restart the NAE-KMIP service, by clicking Restart. This activates the changes on the KMIP interface.
3. Download the certificate.
   a. Navigate to Admin Settings > Interfaces.
   b. Click on the ellipses icon (...) for the KMIP interface and select Edit.
   c. Select Download Current Certificate. The certificate downloads as a text file.
   d. Copy and paste the contents of the downloaded file into a new file with the .pem file type extension.
   This becomes the chained server certificate created in Configuring the storage system (for example, 134006-chained-key_server_9.11.110.162.pem).

Configuring two key servers on the storage system

To complete the configuration and activate the encryption follow step 7 of Configuring the storage system through Activating the encryption.

Thales CipherTrust Manager k170v v2.2.x migration to client certificate authentication key

Migrate the global key management scheme to client certificate authentication key with Thales CipherTrust Manager server k170v 2.2.x on your storage system.

The information in this section:

- Does not include installation of the key server.
- Does not cover all available configuration options. It contains the basic configuration to allow encryption between a FlashSystem A9000R storage system and the CipherTrust Manager key server.
- Uses the storage system management port example IP address `9.11.80.61` and the CipherTrust Manager server address `9.11.110.162`.
  Replace these addresses with the correct IP addresses of your FlashSystem A9000R storage system and of your key server.
- Uses the FlashSystem A9000R machine serial number as a prefix when naming PEM files. For example, if the file is saved as 1340006-ca.pem, *134006* indicates the machine serial number of the storage array. To get the serial number of a FlashSystem A9000R storage array, use the **config_get** CLI command (see Displaying the values of configuration parameters).

Known issues

When using Thales CipherTrust Manager with FlashSystem A9000R, and disabling encryption, it may fail with `ENCRYPT_UNABLE_TO_UPDATE_KEY_DURING_DEACTIVATE_ON_KEYSERVER`. This happens if the system was already encrypt enabled, then disabled and enabled again. This can also happen if the key was in Pre-Active state on the key server when encrypt enabling occurred and if this state was not changed later to 'Active'. In order to prevent this issue from occurring, before encryption, enable the key on the key server side. The key on the key server cannot be deactivated or re-activated if it was already deactivated once. If it is deactivated, a new key will be needed on the key server side.

- **Enabling encryption on the storage system**
  Enable encryption on the storage system side with the CipherTrust Manager global key authentication.
- **Migrating from global key to Client Certificate Authentication**
  Migrate from global key to Client Certificate Authentication.
- **CipherTrust Manager cluster configuration**
  Provide additional protection and maintain the key server cluster by creating a backup key server.
- **Using an external CA certificate**
  Configure an external CA on the CipherTrust Manager and on the storage system.

Enabling encryption on the storage system

Enable encryption on the storage system side with the CipherTrust Manager global key authentication.

Preconfigure the storage system by using the following steps:

1. Verify that the storage system is encryption-ready by running the **state_list** command (see Listing the operational state).
   If encryption is not supported, contact IBM Support.

2. Check if any certificates are defined on the system, using the **pki_list** command (see Listing PKI items).
   Note: Having certificates that are predefined on the system does not cause any issues with the configuration.

Configuring the key server

All procedures that are detailed here take place on the Thales CipherTrust Manager, unless otherwise stated.

1. Create a local Certificate Authority (CA).
   Note:
   - This step might not be necessary if a CA is already in the Local Certificate Authority List If using an External CA, see Using an external CA certificate.
   - This PEM text example and its decoded X.509 are presented for illustrative purposes only and do not completely fit the values shown in the associated screen captures. For example, there may be inconsistencies in the Common Name (CN) and the Certificate Name field values.

   a. Navigate to CA.
   b. From the Local Certificate Authorities section, select New Local CA.
      Important: Use a strong enough key length and algorithm, such as RSA-2048, according to your company's security policies.
   c. Fill in the relevant information and click Create Local CA. The certificate will appear under Pending CAs in the Local Certificate Authority List.
   d. Select the Pending CA and click Self Sign. The CA now appears in the Local Certificate Authority List.
   e. Click on the ellipses icon (...), next to the CA certificate, and select Download.
      This file will be used in later steps to create the client chained certificate. (See step 5 of Configuring the storage system.)

   f. Save the root (local) CA certificate in PEM format. For example, 1340006-ca.pem.
      The PEM format encodes an X.509 certificate.

   g. Check that the root certificate can sign other certificates. Use the following command:

      ```
      # openssl x509 -in 1340006-ca.pem -inform pem -text -noout
      ```

      Note: The following are indications of a root certificate which can sign other certificates:
      - The Issuer and the Subject are the same.
      - There is an X.509 extension indicating it is a CA (`CA:TRUE`).

2. Configure KMIP.
   a. Navigate to Access Management > Registration Tokens.
   b. Click + New Registration Token.
      i. Assign a name for the new token in the Name Prefix field and click Next.
      ii. Select the Local CA that will be used and click Create Token.
      iii. Copy the token and click Done.
      Note: This token will be used in the next step.
   c. Navigate to Admin Setting > Interfaces.
      i. If the KMIP interface does not exist, it needs to be created.
      ii. Click on the ellipses icon (...) and select Edit.
      iii. Configure the following if not already set:
      - Select Enable hard delete and Auto Registration.
      - In the Registration Token field, paste the token that was copied in step 2.b.iii.
      - From the Mode drop-down list, select TLS, verify client cert, allow anonymous logins.
      - From the Local CA for Automatic Server Certificate Generation drop-down list, select your local CA.
      - Add in all local trusted CAs from the Local Trusted CAs drop-down list and click the add (+) button.

      iv. Click Update.
   d. Restart the NAE service.
      i. Navigate to Admin Settings > System Services.
      ii. Reboot the NAE service by clicking Restart.
      Important: This is required to set the CA for KMIP configuration.
3. Retrieve the KMIP server certificate.
   a. Navigate to CA.
      i. Select the newly created Local CA.
      The new kmip.keysecure.local certificate is within the CA.

      ii. Click on the ellipses icon (...) of the kmip.keysecure.local certificate and click Download.
      iii. Save the server certificate in PEM format. For example, key_server_9.11.110.162.pem.

Note: Up to four key servers can be configured at the same time on each FlashSystem A9000R storage array and each one must be configured with its own certificate that is signed by the same CA certificate. Therefore, be sure to identify each server PEM file with the IP address or FQDN of the specific key server for easy identification. When using the **encrypt_keyserver_define** command, be sure to use all existing signed storage array certificates.
The PEM file is used in 6 of Configuring the storage system.

The PEM format encodes an X.509 certificate.

iv. Check that the root certificate can sign other certificates.
Use the following command:

```
# openssl x509 -in key_server_9.11.110.162.pem -inform pem -text -noout
```

Note: The following are indications of a signed certificate which cannot sign other certificates:
- The Issuer and the Subject are *not* the same.
- There is an X.509 extension indicating it is *not* a CA (`CA:FALSE`).

4. Create an AES encryption key.
   a. Navigate to Keys.
   b. To add an AES encryption key, click Create a New Key.
   c. Set the Create Key fields.
      - Fill in the Key Name. The key name is the machine model and serial number of your storage system. For example, `425-1340006`, where *425* is the machine model and *1340006* is the serial number.
      - From the Algorithm drop-down list, select AES.
      - From the Size drop-down list, select 256.

   Click Create.
5. Add the Key Users group.
   a. Set Added Only to off.
   b. In the search groups field, find Key Users.
6. Set Exportable and Deletable to on.
7. On the Default tab, within the KEY ACCESS form, set the Key Owner to global from the drop-down list.
8. On the KMIP tab, navigate to the KMIP | ALIASES & ALTERNATIVE NAMES form.
   a. In the Aliases field, write in the Key Name and click the plus icon (+) to add it.
   b. Click Update.

Configuring the storage system

1. Creating the Certificate Signing Request (CSR) for the KMIP client certificate.
   Create the request on the storage system.
   a. Create the Certificate Signing Request (CSR) on the storage system side.
   This generates a certificate, which is later used to sign with local CA in CipherTrust Manager Server (see step 2 of Configuring the storage system).

   Define the following parameters in the `pki_generate_private_key_and_csr` command (see Generating a private key and CSR):
   - Name is the certificate's symbolic name.
   - Subject is the subject name for the generated certificate request.

   The argument must be formatted as `/type0=value0/type1=value1/type2=...` .
   Include the email of the person who needs to be notified of any certificate or security issues.

   For example,

   ```
   pki_generate_private_key_and_csr name=CipherTrust subject="/C=US/ST=Arizona/L=Tucson/O=IBM/OU=IBM-
   lab/CN=secadmin1/emailAddress=name@company.com"
   CSR
   ```

   where *CipherTrust* is the name and *secadmin1* is the CN (Common Name) username, with the storage array located in Tuscon, Arizona.
   The storage system certificate is displayed.

   b. Copy the entire text of the storage system SSL Certificate Signing Request (CSR), including the `'-----BEGIN CERTIFICATE REQUEST-----' and '-----END CERTIFICATE REQUEST-----'`.
   c. Verify the PKI items, using the **pki_list** command (see Listing PKI items).

   ```
   pki_list
   Name          Fingerprint                          Has signed certificate   Services
   ----------    ---------------------------------- ------------------------ ----------------------------
   1340006       34b52765fbca8c1972931d63e581050c     yes                       XCLI,CIM,IPSEC,KMIP,QUORUM
   CipherTrust    4fd2da6d2ac67e966ae02ad59f0c9147    no                        NONE
   ```

2. Sign the storage system Certificate Signing Request (CSR).
   Use a CipherTrust Manager key secure local CA to sign the Client Certificate Request.
   a. Navigate to CA.
   b. Select a local CA created with the procedures in Creation of a Local CA section.
   c. Select Upload and Sign CSR.
   d. From the Upload Externally Generated CSR form:
      i. Paste the certificate request into the Certificate Request field.
      ii. Select the Client as the Certificate Purpose.
      iii. Specify a Certificate Duration.
   e. Once the form is complete, click Issue Certificate.
3. Download the signed storage system client certificate PEM file.
   a. Click on the ellipses icon (...) next to the newly signed certificate and select Download.
   b. Save the signed client certificate in PEM format. For example: 1340006-client.pem
   The PEM format encodes an X.509 certificate.

4. Check that the client certificate cannot sign other certificates (that it is a leaf certificate).
   Use the following command:

   ```
   # openssl x509 -in 1340006-client.pem -inform pem -text -noout
   ```

   Note: The following are indications of a signed certificate which cannot sign other certificates:
   - The Issuer and the Subject are *not* the same.
   - There is an X.509 extension indicating it is *not* a CA (`CA:FALSE`).
5. Create a chained client certificate PEM file.
   These steps must be done on the Linux host where the two files were downloaded from the key server.

   Use the local root CA PEM file, created and downloaded in step 1 of Configuring the key server (in these examples, 1340006-ca.pem) and the signed storage client certificate PEM file, downloaded in step 3 of Configuring the storage system (in these examples, 1340006-client.pem).

   Ensure that the file used in the CLI command exists on the host where the XCLI Utility is installed. If the files downloaded from CipherTrust Manager are on a different Linux host, the files must be copied to the host where the utility is installed. If needed, a utility, such as SCP (secure copy), may be used to copy the file.

   a. Create a chained client certificate PEM file by combining the local and client certificates.
      Before creating the chained certificate, the certificates can be viewed by using the **cat <certificatename>.pem** command.

      Create the chained certificate using the following command:
      Important: The chained certificate starts with the leaf (for example, the storage system client certificate) and ends with the root certificate (the local CA certificate).
      Verify that the chain is setup in the correct order in order to avoid configuration issues.

      For example:

      ```
      cat 1340006-client.pem 1340006-ca.pem >1340006-chained-client.pem
      ```

   b. Install the chained client certificate on the storage system and update the certificate's related services.
      From the storage system XCLI Utility:
      i. Import the signed PEM certificate, using the **pki_set_pem** command (see Importing a signed certificate).

      ```
      pki_set_pem services=none certificate="$(tr "\n" "*" </root/IBM_Storage_Extended_CLI/1340006-chained-client)"
      ```

      ii. List the PKI items, using the **pki_list** command.
      iii. Remove the KMIP service from the services list of the chained client certificate, using the following `pki_update` CLI command (see Updating a PKI certificate or services).
      The expected output is `Command completed successfully.`

      iv. Add the KMIP service to the CipherTrust certificate, using the **pki_update** command.
      v. Verify that all PKI items are properly updated, using the **pki_list** command.
6. Create a chained server certificate PEM file.
   These steps must be done on the Linux host where the two files were downloaded from the key server.

   Use the local root CA PEM file, created and downloaded in steps 1 of Configuring the key server and the KMIP server certificate (in these examples, 1340006-ca.pem), retrieved in step 3 of Configuring the key server. (in these examples, key_server_9.11.110.162.pem)

   Ensure that the file used in the CLI command exists on the host where the XCLI Utility is installed. If the files downloaded from CipherTrust Manager are on a different Linux host, the files must be copied to the host where the utility is installed. If needed, a utility, such as SCP (secure copy), may be used to copy the file.

   Create a chained server certificate PEM file by combining the local and server certificates. Before creating the chained certificate, the certificates can be viewed by using the **cat <certificatename>.pem** command.

   Create the chained certificate using the following command:
   Important: The chained certificate starts with the leaf (for example, the server certificate) and ends with the root certificate (the local CA certificate).
   Verify that the chain is setup in the correct order in order to avoid configuration issues.

   For example:

   ```
   cat key_server_9.11.110.162.pem 1340006-ca.pem > 1340006-chained-key_server_9.11.110.162.pem
   ```

7. Configure the CipherTrust Manager key server on the storage system.
   On the host that will be used to run CLI commands, make sure that the storage system chained server certificate (for example, 1340006-chained-key_server_9.11.110.162.pem) created in step 6 of Configuring the storage system exists.

   a. Perform the following CLI commands from the storage system:
      Note:
      - Use a Linux machine to perform this step as the storage system expects the certificate data to be streamed.
      - The default port is 5696. This can be changed if necessary.

      i. **encrypt_keyserver_define** with the following parameters defined (see Defining a key server):
         - name: `globalkey`
         - ipv4: the Thales key server
         - master: `yes`
         - keyserver_type: `KeySecure`
         - certificate: the chained server PEM certificate

      For example:

      ```
      encrypt_keyserver_define name=globalkey ipv4=(THALES KeyServer IP) keyserver_type=KeySecure master=yes certificate="$(tr "\n" "*" </root/IBM_Storage_Extended_CLI/1340006-chained-key_server_9.11.110.162.pem)"
      ```

Note: When using the **encrypt_keyserver_define** command, be sure to use all existing signed storage array certificates.

ii. Check the key server, using the **encrypt_keyserver_list** command (see Displaying key server status).
Note: At this point, when encryption on the storage system has still not be enabled, it is okay to see *NOKEY* in the App/Key Status column. This means that there is still no key encrypting the data.

Creating recovery keys

Best practice recommends that recovery keys be configured. When recovery keys are configured, be sure that the command below includes **recovery_keys=yes**.

.For more information regarding recovery keys and how to generate them, see the "Recovery key use and maintenance" section in the *Data-at-rest Encryption for the IBM Spectrum Accelerate Family* Redpaper (REDP-5402).

In order to create recovery keys, at least two Security Admin users must be defined on the storage system.

1. Define at least two Security Admin users on the storage system.
   a. Define each of the Security Admin users, using the **user_define** command (see Defining a new user). For example:

```
user_define user=secadmin1 category=securityadmin password=<password> password_verify=<password>
user_define user=secadmin2 category=securityadmin password=<password> password_verify=<password>
```

   b. Verify the creation of both Security Admins, using the **user_list** command (seeListing users).
2. Create the recovery keys.
   Create two recovery keys, one per Security Admin user.
   a. Define the Security Admin users who are to be receiving the recovery key, using the **encrypt_recovery_key_generate** command (see Generating recovery keys).
   For example:

```
encrypt_recovery_key_generate users=secadmin1,secadmin2  min_req=2
```

   b. Log in with each Security Admin credential and generate the recovery key, using the **encrypt_recovery_key_get** command (see Retrieving the security administrator's recovery key).
   The following is an example of a full recovery key:

```
62807CB1902AM074EDLA4EV8F0C574E40A1564F55570CDEEBED37BC3876789
```

3. Each Security Admin must verify their recovery key, using the **encrypt_recovery_key_verify** command (see Recovering key verification).

Activating the encryption

1. Update the key server, using the **encrypt_enable recovery_keys=yes -y** command (see Enabling encryption).
2. Run the **state_list** command to check that encryption is *Enabled* (see Listing the operational state). If the status is *Enabling*, run the command again, until status is *Enabled*.
3. Check the key server, using the **encrypt_keyserver_list** command (see Displaying key server status).
   The App/Key Status column should show an *Active* status.

Verifying the status

After the encryption keys have been successfully enabled, the flash enclosures and vault devices start changing from a *Ready* to an *Enrolled* state.

- Verify the encryption state of the flash enclosures and vault devices, by running the following CLI commands as a Security Administrator:
  Note: There are corresponding events visible when using the **event_list** command (see Listing events).
  - **flash_enclosure_list -S** (see Listing flash enclosure status)
  - **vault_device_list -S** (see Listing vault devices in the system)

Migrating from global key to Client Certificate Authentication

Migrate from global key to Client Certificate Authentication.

All procedures detailed in this section take place within the Thales CipherTrust Manager Management Console, unless otherwise stated.

Creating a Local User on CipherTrust Manager

If a user has already been created to be used in the Client Certificate Authentication of other devices, skip this section, and proceed to Editing the AES encryption key on CipherTrust Manager.

If a unique user for each device is to be used, this step is mandatory for each and for each FlashSystem A9000R storage array.

If the Certificate Signing Request (CSR) is created with multiple Common Names (CNs), create a user for each CN.

1. Define a user, from Access Management > Users.
   a. Click Create New User.
   b. Define a Username. This username is also used as the Storage Security user.
   c. Fill in the form for the Email, Password, and Certificate Subject Distinguished Name fields.
   d. Select Allow user to login using certificate.
   e. Click Create.
2. Repeat step 1 to create a second user.

3. Configure KMIP.

Editing the AES encryption key on CipherTrust Manager

1. From the Keys menu, select the key that was created in 4 of Configuring the key server, with your storage system model and serial number.
2. In the KEY ACCESS form, use the Key Owner drop-down list to select the user created previously in Creating a Local User on CipherTrust Manager.
   Note: When creating a CSR and using multiple CNs, the key owner is last on the CN list.
3. Scroll to the bottom of the page and select Update.
   The new owner is now reflected on the Keys main page.

Editing the KMIP key server configuration

1. Navigate to Admin Settings > Interfaces.
2. Click on the ellipses icon (...) of the kmip interface and click Edit.
3. From the Mode drop-down list, select TLS, verify client cert, user name taken from client cert, auth request is optional.
4. From the Username Location in Certificate drop-down list, select CN.
5. Click Update.
6. Navigate to Admin Settings > Services.
7. Reboot the NAE service by clicking Restart.

Verifying and configuring the storage system

This section assumes that encryption is enabled with the globalkey user. If this is not the case, refer to Configuring the storage system and enable the global key user before proceeding.

1. Check the CipherTrust Manager Key Server configuration on your storage system.
   To check the configuration use the following CLI commands:
   a. **encrypt_keyserver_check_status** (see Checking key server status).
   b. **encrypt_keyserver_list** (see Displaying key server status).
   If the App/Key Status *does not* show the *Active* status, continue to step 2.

   If the status is *Active*, the migration and configuration is complete. For further protection, create a backup key server (see CipherTrust Manager cluster configuration).

2. Change the Key Status to *Active*.
   a. Navigate to KMIP > Registered Clients.
   b. Click on the ellipses icon (...) for each of the registered clients, and select Delete for each instance.
3. Recheck the Key Status.
   a. Use the commands from step 1.
   The App/Key Status should now be *Active*.

   b. Verify the status again through the storage system events log. For example:

```
2021-03-04 16:09:44 849272 ENCRYPT_UNABLE_TO_RETRIEVE_KEY_FROM_KEYSERVER Failed to retrieve key from key server
'KS' via IPv4 on module 1. Please verify that the key server type and version are supported. If so, please
check its status.
2021-03-04 16:09:44 849273 ENCRYPTION_SKMIP_ERROR           Module 3 reported KS returned error:1 - 1
2021-03-04 16:09:44 849274 ENCRYPT_UNABLE_TO_RETRIEVE_KEY_FROM_KEYSERVER Failed to retrieve key from key server
'KS' via IPv4 on module 3. Please verify that the key server type and version are supported. If so, please
check its status.
2021-03-04 16:09:44 849275 ENCRYPT_UNABLE_TO_RETRIEVE_KEY_FROM_KEYSERVER Failed to retrieve key from key server
'KS' via IPv4 on module 2. Please verify that the key server type and version are supported. If so, please
check its status.
2021-03-04 16:09:44 849276 EMAIL_NOT_SENT                   Sending event
ENCRYPT_UNABLE_TO_RETRIEVE_KEY_FROM_KEYSERVER (849272) to Event Center via us.relay.ibm.com was  waived because
of failed SMTP gateway. It will be not be used until 2021-03-04 16:11:27.
2021-03-04 16:09:44 849277 EMAIL_NOT_SENT                   Sending event ENCRYPTION_SKMIP_ERROR (849273) to
Event Center via us.relay.ibm.com was  waived because of failed SMTP gateway. It will be not be used until
2021-03-04 16:11:27.
2021-03-04 16:09:44 849278 EMAIL_NOT_SENT                   Sending event
ENCRYPT_UNABLE_TO_RETRIEVE_KEY_FROM_KEYSERVER (849274) to Event Center via us.relay.ibm.com was  waived because
of failed SMTP gateway. It will be not be used until 2021-03-04 16:11:27.
2021-03-04 16:09:44 849279 EMAIL_NOT_SENT                   Sending event
ENCRYPT_UNABLE_TO_RETRIEVE_KEY_FROM_KEYSERVER (849275) to Event Center via us.relay.ibm.com was  waived because
of failed SMTP gateway. It will be not be used until 2021-03-04 16:11:27.
2021-03-04 16:10:58 849280 ENCRYPT_KEYSERVER_EDITED         Details of key server '9.11.54.106' were modified.
```

CipherTrust Manager cluster configuration

Provide additional protection and maintain the key server cluster by creating a backup key server.

The backup key server is important in order to avoid a single point of failure (SPOF) in case the key server fails. When a backup key server is in place and a keyserver failure occurs, the backup that is clustered with it can keep running. The preferred practice is that the backup is located in a different physical location than the primary server.

Configuring the cluster

This section assumes that there are two functioning key servers.

Use the CipherTrust Manager Management Console to complete the following procedure.

1. Navigate to Admin Settings > Cluster.
2. Select Manage Cluster > Create Cluster.
3. Configure the local key server and click Save.
4. From the Cluster Configuration page, click Manage Cluster.
5. Click Add Node.
6. Configure the second key server.
    a. From the Public address of the new node field, configure the IP address of the second key server.
    b. Click Add Node. This opens the second CipherTrust Manager key server.
    c. Click Join. This guides you through the process of clustering the two key servers.
    d. Click Finish to complete the process.
    The result displays both keys servers in the cluster in a *Ready* state.

Configuring two key servers on the storage system

This section assumes that one key server is already configured.

Define the second key server, using the **encrypt_keyserver_define** command (see Defining a key server).

Use the same chained server certificate PEM file used in the first key server configuration and ensure that to set `master=no`.
For example:

```
encrypt_keyserver_define name=a9 ipv4=(THALES KeyServer IP) keyserver_type=KeySecure master=no certificate="$(tr "\n" "*"
</root/IBM_Storage_Extended_CLI/chained_key_server_9.11.110.162.pem)"
```

Using an external CA certificate

Configure an external CA on the CipherTrust Manager and on the storage system.

Configuring the external CA on CipherTrust

Be sure that the Certificate Signing Request (CSR) with the external CA is signed with the TLS Web Client Authentication option enabled.

Use the CipherTrust Manager Management Console to complete the following procedure.

1. Upload the external certificate.
    a. From the CA form, select Upload External Certificate.
    b. Paste the external certificate in PEM format.
    c. Click Upload.
2. Add the CA to the trusted list for the KMIP interface.
    a. Navigate to Admin Settings > Interfaces.
    b. Click on the ellipses icon (...) for the kmip interface and select Edit.
    c. From the External Trusted CAs drop-down list, select the external CA that was uploaded in step 1.
    d. Click on the plus icon (+) to add the CA. This adds the CA into the trusted list for the KMIP interface.
    e. Click Update.
    f. Navigate to Admin Settings > Services and select System Restart. This activates the changes on the KMIP interface.
3. Download the certificate.
    a. Navigate to Admin Settings > Interfaces.
    b. Click on the ellipses icon (...) for the kmip interface and select Edit.
    c. Select Download Current Certificate. The certificate downloads as a text file.
    d. Copy and paste the contents of the downloaded file into a new file with the .pem file type extension.
    This becomes the chained server certificate created in Configuring the storage system (for example, 134006-chained-key_server_9.11.110.162.pem).

Configuring two key servers on the storage system

To complete the configuration and activate the encryption follow step 7 of Configuring the storage system through Activating the encryption.

Thales CipherTrust Manager k170v v2.0.0 migration to client certificate authentication key

Migrate the global key management scheme to client certificate authentication key with Thales CipherTrust Manager server k170v 2.0.0 on your storage system.

The information in this section:

- Does not include installation of the key server.
- Does not cover all available configuration options. It contains the basic configuration to allow encryption between a FlashSystem A9000R storage system and the CipherTrust Manager key server.
- Uses the storage system management port example IP address `9.11.80.61` and the CipherTrust Manager server address `9.11.110.162`.

Replace these addresses with the correct IP addresses of your FlashSystem A9000R storage system and of your key server.

- Uses the FlashSystem A9000R machine serial number as a prefix when naming PEM files. For example, if the file is saved as 1340006-ca.pem, *134006* indicates the machine serial number of the storage array. To get the serial number of a FlashSystem A9000R storage array, use the **config_get** CLI command (see Displaying the values of configuration parameters).

Known issues

When using Thales CipherTrust Manager with FlashSystem A9000R, and disabling encryption, it may fail with `ENCRYPT_UNABLE_TO_UPDATE_KEY_DURING_DEACTIVATE_ON_KEYSERVER`. This happens if the system was already encrypt enabled, then disabled and enabled again. This can also happen if the key was in Pre-Active state on the key server when encrypt enabling occurred and if this state was not changed later to 'Active'. In order to prevent this issue from occurring, before encryption, enable the key on the key server side. The key on the key server cannot be deactivated or re-activated if it was already deactivated once. If it is deactivated, a new key will be needed on the key server side.

- **Enabling encryption on the storage system**
  Enable encryption on the storage system side with the CipherTrust Manager global key authentication.
- **Migrating from global key to Client Certificate Authentication**
  Migrate from global key to Client Certificate Authentication.
- **CipherTrust Manager cluster configuration**
  Provide additional protection and maintain the key server cluster by creating a backup key server.

Enabling encryption on the storage system

Enable encryption on the storage system side with the CipherTrust Manager global key authentication.

Preconfigure the storage system by using the following steps:

1. Verify that the storage system is encryption-ready by running the **state_list** command (see Listing the operational state).
   If encryption is not supported, contact IBM Support.
2. Check if any certificates are defined on the system, using the **pki_list** command (see Listing PKI items).
   Note: Having certificates that are predefined on the system does not cause any issues with the configuration.

Configuring the key server

All procedures that are detailed here take place on the Thales CipherTrust Manager, unless otherwise stated.

1. Create a local Certificate Authority (CA).
   Note:
   - This step might not be necessary if a CA is already in the Local Certificate Authority List.
   - This PEM text example and its decoded X.509 are presented for illustrative purposes only and do not completely fit the values shown in the associated screen captures. For example, there may be inconsistencies in the Common Name (CN) and the Certificate Name field values.

   a. Navigate to Keys & Access Management > CA.
   b. From the Local Certificate Authorities section, select New Local CA.
      Important: Use a strong enough key length and algorithm, such as RSA-2048, according to your company's security policies.
   c. Fill in the relevant information and click Create Local CA. The certificate will appear under Pending CAs in the Local Certificate Authority List.
   d. Select the Pending CA and click Self Sign. The CA now appears in the Local Certificate Authority List.
   e. Click on the ellipses icon (...), next to the CA certificate, and select Download.
      This file will be used in later steps to create the client chained certificate. (See step 5 of Configuring the storage systemc.)

   f. Save the root (local) CA certificate in PEM format. For example, 1340006-ca.pem.
      The PEM format encodes an X.509 certificate.

   g. Check that the root certificate can sign other certificates. Use the following command:

      ```
      # openssl x509 -in 1340006-ca.pem -inform pem -text -noout
      ```

      Note: The following are indications of a root certificate which can sign other certificates:
      - The Issuer and the Subject are the same.
      - There is an X.509 extension indicating it is a CA (**`CA:TRUE`**).
2. Configure KMIP.
   a. Navigate to Keys & Access Management > Registration Tokens.
   b. Click + New Registration Token.
      i. Assign a name for the new token in the Name Prefix field and click Next.
      ii. Select the Local CA that will be used and create the token..
      iii. Copy the token and click Done.
      Note: This token will be used in the next step.
   c. Navigate to Admin Setting > System > Interfaces.
      i. If the KMIP interface does not exist, it needs to be created.
      ii. Click on the ellipses icon (...) and select Edit.
      iii. Configure the following, if not already set:
      - Select Enable hard delete and Auto Registration.
      - In the Registration Token field, paste the token that was copied in step 2.b.iii.
      - From the Mode drop-down list, select TLS, verify client cert, allow anonymous logins.
      - From the Local CA for Automatic Server Certificate Generation drop-down list, select your local CA.

- Add in all local trusted CAs from the Local Trusted CAs drop-down list and click the add (+) button.

iv. Click Update.

v. If required, restart the KMIP interface service.

1. Navigate to Admin Setting > System > Services.
2. Click Restart on the KMIP interface and confirm the restart.
3. Return to Admin Setting > System > Interfaces and select the KMIP interface.
   Changes are saved automatically.
4. Navigate to Admin Setting > System > Services.
5. Reboot the CipherTrust Manager server by clicking on System Restart.
   Important: This is required to set the CA for KMIP configuration.

3. Retrieve the KMIP server certificate.

a. Navigate to Keys & Access Management > CA.

i. Select the newly created Local CA.
The new kmip.keysecure.local certificate is within the CA.

ii. Click on the ellipses icon (...) of the kmip.keysecure.local certificate and click Download.

iii. Save the server certificate in PEM format. For example, key_server_9.11.110.162.pem.
Note: Up to four key servers can be configured at the same time on each FlashSystem A9000R storage array and each one must be configured with its own certificate that is signed by the same CA certificate. Therefore, be sure to identify each server PEM file with the IP address or FQDN of the specific key server for easy identification. When using the **encrypt_keyserver_define** command, be sure to use all existing signed storage array certificates.
The PEM file is used in 6 of Configuring the storage system.

The PEM format encodes an X.509 certificate.

iv. Check that the root certificate can sign other certificates.
Use the following command:

```
# openssl x509 -in key_server_9.11.110.162.pem -inform pem -text -noout
```

Note: The following are indications of a signed certificate which cannot sign other certificates:
- The Issuer and the Subject are *not* the same.
- There is an X.509 extension indicating it is *not* a CA (`CA:FALSE`).

4. Create an AES encryption key.

a. Navigate to Keys.

b. To add an AES encryption key, click Create a New Key.

c. Set the Create Key fields.
- Fill in the Key Name. The key name is the machine model and serial number of your storage system. For example, `425-1340006`, where *425* is the machine model and *1340006* is the serial number.
- Set the Key Owner to *global*.
- From the Algorithm drop-down list, select AES.
- From the Size drop-down list, select 256.
- Add the Key Users group.
- From the Key Properties form in the Alias (KMIP Name) field add in the key name again. Click the plus icon (+) to add it.
- Clear the Prevent this key from being deleted and Prevent this key from being exported checkboxes.

Click Create.

Configuring the storage system

1. Creating the Certificate Signing Request (CSR) for the KMIP client certificate.
Create the request on the storage system.

a. Create the Certificate Signing Request (CSR) on the storage system side.
This generates a certificate, which is later used to sign with local CA in CipherTrust Manager Server (see step 2 of Configuring the storage system).

Define the following parameters in the **pki_generate_private_key_and_csr** command (see Generating a private key and CSR):
- Name is the certificate's symbolic name.
- Subject is the subject name for the generated certificate request.

The argument must be formatted as **/type0=value0/type1=value1/type2=...** .
Include the email of the person who needs to be notified of any certificate or security issues.

For example,

```
pki_generate_private_key_and_csr name=CipherTrust subject="/C=US/ST=Arizona/L=Tucson/O=IBM/OU=IBM-
lab/CN=secadmin1/emailAddress=name@company.com"
CSR
```

where *CipherTrust* is the name and *secadmin1* is the CN (Common Name) username, with the storage array located in Tuscon, Arizona.
The storage system certificate is displayed.

b. Copy the entire text of the storage system SSL Certificate Signing Request (CSR), including the **'-----BEGIN CERTIFICATE REQUEST-----' and '-----END CERTIFICATE REQUEST-----'**.

c. Verify the PKI items, using the **pki_list** command (see Listing PKI items).

```
pki_list
Name          Fingerprint                          Has signed certificate   Services
----------    ---------------------------------- ------------------------ ----------------------------
1340006       34b52765fbca8c1972931d63e581050c     yes                       XCLI,CIM,IPSEC,KMIP,QUORUM
CipherTrust    4fd2da6d2ac67e966ae02ad59f0c9147    no                        NONE
```

2. Sign the storage system Certificate Signing Request (CSR).

Use a CipherTrust Manager key secure local CA to sign the Client Certificate Request.

a. Navigate to Keys & Access Management > CA.
b. Select a local CA created with the procedures in Creation of a Local CA section.
c. Select Upload and Sign CSR.
d. From the Upload Externally Generated CSR form:
   i. Paste the certificate request into the Certificate Request field.
   ii. Select the Client as the Certificate Purpose.
   iii. Specify a Certificate Duration.
e. Once the form is complete, click Issue Certificate.

3. Download the signed storage system client certificate PEM file.
   a. Click on the ellipses icon (...) next to the newly signed certificate and select Download.
   b. Save the signed client certificate in PEM format. For example: 1340006-client.pem
      The PEM format encodes an X.509 certificate.

4. Check that the client certificate cannot sign other certificates (that it is a leaf certificate).
   Use the following command:

```
# openssl x509 -in 1340006-client.pem -inform pem -text -noout
```

   Note: The following are indications of a signed certificate which cannot sign other certificates:
   - The Issuer and the Subject are *not* the same.
   - There is an X.509 extension indicating it is *not* a CA (**`CA:FALSE`**).

5. Create a chained client certificate PEM file.
   These steps must be done on the Linux host where the two files were downloaded from the key server.

   Use the local root CA PEM file, created and downloaded in step 1 of Configuring the key server (in these examples, 1340006-ca.pem) and the signed storage client certificate PEM file, downloaded in step 3 of Configuring the storage system (in these examples, 1340006-client.pem).

   Ensure that the file used in the CLI command exists on the host where the XCLI Utility is installed. If the files downloaded from CipherTrust Manager are on a different Linux host, the files must be copied to the host where the utility is installed. If needed, a utility, such as SCP (secure copy), may be used to copy the file.

   a. Create a chained client certificate PEM file by combining the local and client certificates.
      Before creating the chained certificate, the certificates can be viewed by using the **cat <certificatename>.pem** command.

      Create the chained certificate using the following command:
      Important: The chained certificate starts with the leaf (for example, the storage system client certificate) and ends with the root certificate (the local CA certificate).
      Verify that the chain is setup in the correct order in order to avoid configuration issues.

      For example:

```
cat 1340006-client.pem 1340006-ca.pem >1340006-chained-client.pem
```

   b. Install the chained client certificate on the storage system and update the certificate's related services.
      From the storage system XCLI Utility:
      i. Import the signed PEM certificate, using the **pki_set_pem** command (see Importing a signed certificate).

```
pki_set_pem services=none certificate="$(tr "\n" "*" </root/IBM_Storage_Extended_CLI/1340006-chained-client)"
```

      ii. List the PKI items, using the **pki_list** command.
      iii. Remove the KMIP service from the services list of the chained client certificate, using the following **`pki_update`** CLI command (see Updating a PKI certificate or services).
      The expected output is `Command completed successfully.`

      iv. Add the KMIP service to the CipherTrust certificate, using the **pki_update** command.
      v. Verify that all PKI items are properly updated, using the **pki_list** command.

6. Create a chained server certificate PEM file.
   These steps must be done on the Linux host where the two files were downloaded from the key server.

   Use the local root CA PEM file, created and downloaded in step 1 of Configuring the key server and the KMIP server certificate (in these examples, 1340006-ca.pem), retrieved in step 3 of Configuring the key server. (in these examples, key_server_9.11.110.162.pem)

   Ensure that the file used in the CLI command exists on the host where the XCLI Utility is installed. If the files downloaded from CipherTrust Manager are on a different Linux host, the files must be copied to the host where the utility is installed. If needed, a utility, such as SCP (secure copy), may be used to copy the file.

   Create a chained server certificate PEM file by combining the local and server certificates. Before creating the chained certificate, the certificates can be viewed by using the **cat <certificatename>.pem** command.

   Create the chained certificate using the following command:
   Important: The chained certificate starts with the leaf (for example, the server certificate) and ends with the root certificate (the local CA certificate).
   Verify that the chain is setup in the correct order in order to avoid configuration issues.

   For example:

```
cat key_server_9.11.110.162.pem 1340006-ca.pem > 1340006-chained-key_server_9.11.110.162.pem
```

7. Configure the CipherTrust Manager key server on the storage system.
   On the host that will be used to run CLI commands, make sure that the storage system chained server certificate (for example, 1340006-chained-key_server_9.11.110.162.pem) created in step 6 of Configuring the storage system exists.

a. Perform the following CLI commands from the storage system:
Note:
- Use a Linux machine to perform this step as the storage system expects the certificate data to be streamed.
- The default port is 5696. This can be changed if necessary.

i. **encrypt_keyserver_define** with the following parameters defined (see Defining a key server):
- name: `globalkey`
- ipv4: the Thales key server
- master: `yes`
- keyserver_type: `KeySecure`
- certificate: the chained server PEM certificate

For example:

```
encrypt_keyserver_define name=globalkey ipv4=(THALES KeyServer IP) keyserver_type=KeySecure master=yes
certificate="$(tr "\n" "*" </root/IBM_Storage_Extended_CLI/1340006-chained-key_server_9.11.110.162.pem)"
```

ii. Check the key server, using the **encrypt_keyserver_list** command (see Displaying key server status).
Note: At this point, when encryption on the storage system has still not be enabled, it is okay to see *NOKEY* in the App/Key Status column. This means that there is still no key encrypting the data.

Creating recovery keys

Best practice recommends that recovery keys be configured. When recovery keys are configured, be sure that the command below includes **recovery_keys=yes**.

.For more information regarding recovery keys and how to generate them, see the "Recovery key use and maintenance" section in the *Data-at-rest Encryption for the IBM Spectrum Accelerate Family* Redpaper (REDP-5402).

In order to create recovery keys, at least two Security Admin users must be defined on the storage system.

1. Define at least two Security Admin users on the storage system.
   a. Define each of the Security Admin users, using the **user_define** command (see Defining a new user). For example:

```
user_define user=secadmin1 category=securityadmin password=<password> password_verify=<password>
user_define user=secadmin2 category=securityadmin password=<password> password_verify=<password>
```

   b. Verify the creation of both Security Admins, using the **user_list** command (seeListing users).
2. Create the recovery keys.
   Create two recovery keys, one per Security Admin user.
   a. Define the Security Admin users who are to be receiving the recovery key, using the **encrypt_recovery_key_generate** command (see Generating recovery keys).
   For example:

```
encrypt_recovery_key_generate users=secadmin1,secadmin2  min_req=2
```

   b. Log in with each Security Admin credential and generate the recovery key, using the **encrypt_recovery_key_get** command (see Retrieving the security administrator's recovery key).
   The following is an example of a full recovery key:

```
62807CB1902AM074EDLA4EV8F0C574E40A1564F55570CDEEBED37BC3876789
```

3. Each Security Admin must verify their recovery key, using the **encrypt_recovery_key_verify** command (see Recovering key verification).

Activating the encryption

1. Update the key server, using the **encrypt_enable recovery_keys=yes -y** command (see Enabling encryption).
2. Run the **state_list** command to check that encryption is *Enabled* (see Listing the operational state). If the status is *Enabling*, run the command again, until status is *Enabled*.
3. Check the key server, using the **encrypt_keyserver_list** command (see Displaying key server status).
   The App/Key Status column should show an *Active* status.

Verifying the status

After the encryption keys have been successfully enabled, the flash enclosures and vault devices start changing from a *Ready* to an *Enrolled* state.

- Verify the encryption state of the flash enclosures and vault devices, by running the following CLI commands as a Security Administrator:
  Note: There are corresponding events visible when using the **event_list** command (see Listing events).
  - **flash_enclosure_list -S** (see Listing flash enclosure status)
  - **vault_device_list -S** (see Listing vault devices in the system)

Migrating from global key to Client Certificate Authentication

Migrate from global key to Client Certificate Authentication.

All procedures detailed in this section take place within the Thales CipherTrust Manager Management Console, unless otherwise stated.

Creating a Local User on CipherTrust Manager

If a user has already been created to be used in the Client Certificate Authentication of other devices, skip this section, and proceed to Editing the AES encryption key on CipherTrust Manager.

If a unique user for each device is to be used, this step is mandatory for each and for each FlashSystem A9000R storage array.

If the Certificate Signing Request (CSR) is created with multiple Common Names (CNs), create a user for each CN.

1. Define a user, from Keys & Access Management > Users.
   a. Click Create New User.
   b. Define a Username. This username is also used as the Storage Security user.
   c. Fill in the form for the Email, Password, and Certificate Subject Distinguished Name fields.
   d. Select Allow user to login using certificate.
   e. Click Create.
2. Repeat step 1 to create a second user.

Editing the AES encryption key on CipherTrust Manager

1. From the Keys & Access Management > Keys menu, select the key that was created in 4 of Configuring the key server, with your storage system model and serial number.
2. Click on the ellipses icon (...) and select Edit.
3. In the Key Metadata form, use the Key Owner drop-down list to select the user created previously in Creating a Local User on CipherTrust Manager.
   Note: When creating a CSR and using multiple CNs, the key owner is last on the CN list.
4. Scroll to the bottom of the page and select Update.
   The new owner is now reflected on the Keys main page.

Editing the KMIP key server configuration

1. Navigate to Admin Settings > Interfaces.
2. Click on the ellipses icon (...) of the kmip interface and click Edit.
3. From the Mode drop-down list, select TLS, verify client cert, user name taken from client cert, auth request is optional.
4. From the Username Location in Certificate drop-down list, select CN.
5. Click Update.
6. From Admin Settings > System > Services restart the KMIP service, clicking Restart.

Verifying and configuring the storage system

This section assumes that encryption is enabled with the globalkey user. If this is not the case, refer to Configuring the storage system and enable the global key user before proceeding.

1. Check the CipherTrust Manager Key Server configuration on your storage system.
   To check the configuration use the following CLI commands:
   a. **encrypt_keyserver_check_status** (see Checking key server status).
   b. **encrypt_keyserver_list** (see Displaying key server status).
   If the App/Key Status *does not* show the *Active* status, continue to step 2.

   If the status is *Active*, the migration and configuration is complete. For further protection, create a backup key server (see CipherTrust Manager cluster configuration).

2. Change the Key Status to *Active*.
   a. Navigate to KMIP > Registered Clients.
   b. Click on the ellipses icon (...) for each of the registered clients, and select Delete for each instance.
   c. From Admin Settings > System > Services restart the KMIP service, clicking Restart.
3. Recheck the Key Status.
   a. Use the commands from step 1.
   The App/Key Status should now be *Active*.

   b. Verify the status again through the storage system events log. For example:

```
2021-03-04 16:09:44 849272 ENCRYPT_UNABLE_TO_RETRIEVE_KEY_FROM_KEYSERVER Failed to retrieve key from key server
'KS' via IPv4 on module 1. Please verify that the key server type and version are supported. If so, please
check its status.
2021-03-04 16:09:44 849273 ENCRYPTION_SKMIP_ERROR           Module 3 reported KS returned error:1 - 1
2021-03-04 16:09:44 849274 ENCRYPT_UNABLE_TO_RETRIEVE_KEY_FROM_KEYSERVER Failed to retrieve key from key server
'KS' via IPv4 on module 3. Please verify that the key server type and version are supported. If so, please
check its status.
2021-03-04 16:09:44 849275 ENCRYPT_UNABLE_TO_RETRIEVE_KEY_FROM_KEYSERVER Failed to retrieve key from key server
'KS' via IPv4 on module 2. Please verify that the key server type and version are supported. If so, please
check its status.
2021-03-04 16:09:44 849276 EMAIL_NOT_SENT                   Sending event
ENCRYPT_UNABLE_TO_RETRIEVE_KEY_FROM_KEYSERVER (849272) to Event Center via us.relay.ibm.com was  waived because
of failed SMTP gateway. It will be not be used until 2021-03-04 16:11:27.
2021-03-04 16:09:44 849277 EMAIL_NOT_SENT                   Sending event ENCRYPTION_SKMIP_ERROR (849273) to
Event Center via us.relay.ibm.com was  waived because of failed SMTP gateway. It will be not be used until
2021-03-04 16:11:27.
2021-03-04 16:09:44 849278 EMAIL_NOT_SENT                   Sending event
ENCRYPT_UNABLE_TO_RETRIEVE_KEY_FROM_KEYSERVER (849274) to Event Center via us.relay.ibm.com was  waived because
of failed SMTP gateway. It will be not be used until 2021-03-04 16:11:27.
2021-03-04 16:09:44 849279 EMAIL_NOT_SENT                   Sending event
ENCRYPT_UNABLE_TO_RETRIEVE_KEY_FROM_KEYSERVER (849275) to Event Center via us.relay.ibm.com was  waived because
```

```
of failed SMTP gateway. It will be not be used until 2021-03-04 16:11:27.
2021-03-04 16:10:58 849280 ENCRYPT_KEYSERVER_EDITED         Details of key server '9.11.54.106' were modified.
```

CipherTrust Manager cluster configuration

Provide additional protection and maintain the key server cluster by creating a backup key server.

The backup key server is important in order to avoid a single point of failure (SPOF) in case the key server fails. When a backup key server is in place and a keyserver failure occurs, the backup that is clustered with it can keep running. The preferred practice is that the backup is located in a different physical location than the primary server.

Configuring the cluster

This section assumes that there are two functioning key servers.

Use the CipherTrust Manager Management Console to complete the following procedure.

1. Navigate to Admin Settings > System > Cluster.
2. Select Manage Cluster > Add Node.
3. Configure the local key server and click Save.
4. From the Cluster Configuration page, click Manage Cluster.
5. Click Add Node.
6. Configure the second key server.
    a. From the Public address of the new node field, configure the IP address of the second key server.
    b. Click Add Node. This opens the second CipherTrust Manager key server.
    c. Click Join. This guides you through the process of clustering the two key servers.
    d. Click Finish to complete the process.
    The result displays both keys servers in the cluster in a *Ready* state.

Configuring two key servers on the storage system

This section assumes that one key server is already configured.

Define the second key server, using the **encrypt_keyserver_define** command (see Defining a key server).

Use the same chained server certificate PEM file used in the first key server configuration and ensure that to set `master=no`.
For example:

```
encrypt_keyserver_define name=a9 ipv4=(THALES KeyServer IP) keyserver_type=KeySecure master=no certificate="$(tr "\n" "*"
</root/IBM_Storage_Extended_CLI/chained_key_server_9.11.110.162.pem)"
```

Thales SafeNet KeySecure 8.11.1 migration to client certificate authentication key

Migrate the global key management scheme to client certificate authentication key with Thales SafeNet KeySecure 8.11.1 on your storage system.

The information in this section:

- Does not include installation of the key server.
- Does not cover all available configuration options. It contains the basic configuration to allow encryption between a FlashSystem A9000R storage system and the SafeNet KeySecure key server.
- Uses the storage system management port example IP address `9.11.80.61` and the SafeNet KeySecure server address `9.11.110.162`.
  Replace these addresses with the correct IP addresses of your FlashSystem A9000R storage system and of your key server.
- Uses the FlashSystem A9000R machine serial number as a prefix when naming PEM files. For example, if the file is saved as 1340006-ca.pem, *134006* indicates the machine serial number of the storage array. To get the serial number of a FlashSystem A9000R storage array, use the **config_get** CLI command (see Displaying the values of configuration parameters).

Known issues

When using SafeNet KeySecure key server with FlashSystem A9000R, and disabling encryption, it may fail with `ENCRYPT_UNABLE_TO_UPDATE_KEY_DURING_DEACTIVATE_ON_KEYSERVER`. This happens if the system was already encrypt enabled, then disabled and enabled again. This can also happen if the key was in Pre-Active state on the key server when encrypt enabling occurred and if this state was not changed later to 'Active'. In order to prevent this issue from occurring, before encryption, enable the key on the key server side. The key on the key server cannot be deactivated or re-activated if it was already deactivated once. If it is deactivated, a new key will be needed on the key server side.

- **Enabling encryption on the storage system**
  Enable encryption on the storage system side with the SafeNet KeySecure global key authentication.
- **Migrating from global key to Client Certificate Authentication**
  Migrate from global key to Client Certificate Authentication.

Enabling encryption on the storage system

Enable encryption on the storage system side with the SafeNet KeySecure global key authentication.

Preconfigure the storage system by using the following steps:

1. Verify that the storage system is encryption-ready by running the **state_list** command (see Listing the operational state).
   If encryption is not supported, contact IBM Support.

2. Check if any certificates are defined on the system, using the **pki_list** command (see Listing PKI items).
   Note: Having certificates that are predefined on the system does not cause any issues with the configuration.

Configuring the key server

All procedures that are detailed here take place on the Thales SafeNet KeySecure key server, unless otherwise stated.

1. Configure the High Security settings.
   Navigate to Security > High Security and verify that Disable Creation and Use of Global Keys is **not** selected.

2. Create a local Certificate Authority (CA).
   Note:
   - This step might not be necessary if a CA is already in the Local Certificate Authority List.
   - This PEM text example and its decoded X.509 are presented for illustrative purposes only and do not completely fit the values shown in the associated screen captures. For example, there may be inconsistencies in the Common Name (CN) and the Certificate Name field values.

   a. Navigate to Security > Local CAs.
   b. From the Local Certificate Authorities section, populate all of the blank fields.
      Important: Use a strong enough key length and algorithm, such as RSA-2048, according to your company's security policies.
   c. Make sure that Self-signed Root CA is selected.
   d. Click Create.
      The Local Certificate Authority List appears.

   e. Select the CA certificate and click Download.
      This file will be used in later steps to create the client chained certificate. (See step 5 of Configuring the storage system.)

   f. Save the root (local) CA certificate in PEM format. For example, 1340006-ca.pem.
      The PEM format encodes an X.509 certificate.

   g. Check that the root certificate can sign other certificates. Use the following command:

      ```
      # openssl x509 -in 1340006-ca.pem -inform pem -text -noout
      ```

      Note: The following are indications of a root certificate which can sign other certificates:
      - The Issuer and the Subject are the same.
      - There is an X.509 extension indicating it is a CA (`CA:TRUE`).

3. Add the local CA to the KeySecure's trusted CA list.
   a. Navigate to Security > Trusted CA Lists.
   b. From the Trusted Certificate Authority List Profiles form, select an existing profile for KMIP connections or create a new one.
      - Select the existing Default profile and click Properties.
      - Add a specialized profile for KMIP connections.
        Create the new profile by clicking Add. Enter a name for the profile and click Save. Select the profile and click Properties.

   c. From the Trusted Certificate Authority List, click Edit.
      The Trusted CAs and Available CAs lists are displayed.

   d. Select the local certificate authority and move it to the Trusted CAs list, by clicking Add.
   e. Click Save.

4. Create an SSL Certificate for the KMIP key server and sign it with the local CA.
   Note: This step might not be necessary if the customer environment already has an SSL Server Certificate configured for the KeySecure key server.
   a. Create an SSL Certificate for the KMIP key server and sign it with the local CA.
      i. Navigate to Security > SSL Certificates.
      ii. From the Create Certificate Request pane, populate all of the blank fields.
      iii. Click Create Certificate Request.
   b. Verify that the certificate was added from the Certificate List. Select the SSL certificate and click Properties.
   c. Copy the entire text of the SSL certificate, including the `-----BEGIN CERTIFICATE REQUEST-----` and `-----END CERTIFICATE REQUEST-----` text.
   d. Go to Security > Local CAs and click Sign Request.
      The Sign Certificate Request form opens.

   e. Paste the SSL certificate text, copied in step 4.c into the Certificate Request box.
   f. From the Certificate Purpose options, select Server and click Sign Request.
   g. From the CA Certificate Information form, copy the entire certificate text, including the `-----BEGIN CERTIFICATE REQUEST-----` and `-----END CERTIFICATE REQUEST-----` text.
   h. Go to Security > SSL Certificates and click Properties.
      The Certificate Request Information page opens.

   i. At the bottom of the page, click Install Certificate.
   j. Paste the certificate text copied in step 4.g and click Save.

k. From Security > SSL Certificates > Certificate and CA Configuration verify that the Certificate Status in the Certificate List is *Active*.
l. Click Properties.
m. Save the server certificate in PEM format by clicking Download from the Certificate Information form.
n. Save the server certificate in PEM format. For example, key_server_9.11.110.162.pem.
Note: Up to four key servers can be configured at the same time on each FlashSystem A9000R storage array and each one must be configured with its own certificate that is signed by the same CA certificate. Therefore, be sure to identify each server PEM file with the IP address or FQDN of the specific key server for easy identification. When using the **encrypt_keyserver_define** command, be sure to use all existing signed storage array certificates.
The PEM file is used in 6 of Configuring the storage system.

The PEM format encodes an X.509 certificate.

o. Check that the root certificate can sign other certificates.
Use the following command:

```
# openssl x509 -in key_server_9.11.110.162.pem -inform pem -text -noout
```

Note: The following are indications of a signed certificate which cannot sign other certificates:
- The Issuer and the Subject are *not* the same.
- There is an X.509 extension indicating it is *not* a CA (**CA:FALSE**).

5. Add a KMIP key server.
a. Navigate to Device > Key Server and click Add.
Note: Only the KMIP server should remain. This is a strong security recommendation.
b. Select and enter the following properties within the Cryptographic Key Server Properties form:
- Protocol: *KMIP*
- Port: 5696
- Select Use SSL.
- Select the Server Certificate defined in step 4 of Configuring the key server from the drop-down list.

c. Click Save.
d. Select the KMIP protocol and click Properties.
The Cryptographic Key Server Configuration page appears.

e. From the Authentication Settings from, click Edit.
f. From KeySecure > Device tab > Key Server > KMIP edit the KMIP key server Authentication Settings.
- Change Client Certificate Authentication from *Not Used* to *used for SSL session only*.
- Select the correct Trusted CA List Profile from the drop-down list.
- Select NONE from the Username Field in Client Certificate drop-down list.

Click Save.

6. Create an AES encryption key.
a. Navigate to Security > Keys > Key List.
b. Click Create Key to add an AES encryption key.
c. Set the Create Key fields.
- Fill in the Key Name. The key name is the machine model and serial number of your storage system. For example, 425-1340006, where *425* is the machine model and *1340006* is the serial number.
- Leave the Owner Username field blank.
- From the Algorithm drop-down list, select AES-256.
- Select Deletable and Exportable.

Click Create.

Configuring the storage system

1. Creating the Certificate Signing Request (CSR) for the KMIP client certificate.
Create the request on the storage system.
a. Create the Certificate Signing Request (CSR) on the storage system side.
This generates a certificate, which is later used to sign with local CA in the SafeNet KeySecure server (see step 2 of Configuring the storage system).

Define the following parameters in the **pki_generate_private_key_and_csr** command (see Generating a private key and CSR):
- Name is the certificate's symbolic name.
- Subject is the subject name for the generated certificate request.

The argument must be formatted as **/type0=value0/type1=value1/type2=...** .
Include the email of the person who needs to be notified of any certificate or security issues.

For example,

```
pki_generate_private_key_and_csr name=csrforks subject="/C=US/ST=Arizona/L=Tucson/O=IBM/OU=IBM-
lab/CN=secadmin1/emailAddress=name@company.com"
CSR
```

where *csrforks* is the name and *secadmin1* is the CN (Common Name) username, with the storage array located in Tuscon, Arizona.
The storage system certificate is displayed.

b. Copy the entire text of the storage system SSL Certificate Signing Request (CSR), including the **'-----BEGIN CERTIFICATE REQUEST-----' and '-----END CERTIFICATE REQUEST-----'**.
c. Verify the PKI items, using the **pki_list** command (see Listing PKI items).

```
pki_list
Name         Fingerprint                          Has signed certificate   Services
----------   ---------------------------------- ------------------------ ----------------------------
```

```
1340006     34b52765fbca8c1972931d63e581050c   yes                    XCLI,CIM,IPSEC,KMIP,QUORUM
csrforks    4fd2da6d2ac67e966ae02ad59f0c9147   no                     NONE
```

2. Sign the storage system Certificate Signing Request (CSR).
   Use a KeySecure key secure local CA to sign the Client Certificate Request.
   a. Navigate to Security > Local CAs > Local Certificate Authority List.
   b. Select a local CA created in step 2 of Configuring the key server and click Sign Request.
   c. On the Sign Certificate Request form.
      i. Paste the certificate request copied in step 1.b into the Certificate Request.
      ii. From the Certificate Purpose options, select Client.
      iii. Specify a Certificate Duration.
   Click Sign Request.
3. Download the signed storage system client certificate PEM file.
   a. Click Download to download the certificate.
   b. Save the signed client certificate in PEM format. For example: 1340006-client.pem
   The PEM format encodes an X.509 certificate.

4. Check that the client certificate cannot sign other certificates (that it is a leaf certificate).
   Use the following command:

```
# openssl x509 -in 1340006-client.pem -inform pem -text -noout
```

   Note: The following are indications of a signed certificate which cannot sign other certificates:
   - The Issuer and the Subject are *not* the same.
   - There is an X.509 extension indicating it is *not* a CA (`CA:FALSE`).
5. Create a chained client certificate PEM file.
   These steps must be done on the Linux host where the two files were downloaded from the key server.

   Use the local root CA PEM file, created and downloaded in step 2 of Configuring the key server (in these examples, 1340006-ca.pem) and the signed storage client certificate PEM file, downloaded in step 3 of Configuring the storage system (in these examples, 1340006-client.pem).

   Ensure that the file used in the CLI command exists on the host where the XCLI Utility is installed. If the files downloaded from SafeNet KeySecure are on a different Linux host, the files must be copied to the host where the utility is installed. If needed, a utility, such as SCP (secure copy), may be used to copy the file.

   a. Create a chained client certificate PEM file by combining the local and client certificates.
      Before creating the chained certificate, the certificates can be viewed by using the **cat <certificatename>.pem** command.

      Create the chained certificate using the following command:
      Important: The chained certificate starts with the leaf (for example, the storage system client certificate) and ends with the root certificate (the local CA certificate).
      Verify that the chain is setup in the correct order in order to avoid configuration issues.

      For example:

```
cat 1340006-client.pem 1340006-ca.pem >1340006-chained-client.pem
```

   b. Install the chained client certificate on the storage system and update the certificate's related services.
      From the storage system XCLI Utility:
      i. Import the signed PEM certificate, using the **pki_set_pem** command (see Importing a signed certificate).

```
pki_set_pem services=none certificate="$(tr "\n" "*" </root/IBM_Storage_Extended_CLI/1340006-chained-client)"
```

      ii. List the PKI items, using the **pki_list** command.
      iii. Remove the KMIP service from the services list of the chained client certificate, using the following `pki_update` CLI command (see Updating a PKI certificate or services).
      The expected output is `Command completed successfully.`

      iv. Add the KMIP service to the csrforks, using the **pki_update** command.

```
pki_update name=csrforks services=KMIP
```

      v. Verify that all PKI items are properly updated, using the **pki_list** command.
6. Create a chained server certificate PEM file.
   These steps must be done on the Linux host where the two files were downloaded from the key server.

   Use the local root CA PEM file, created and downloaded in step 2 of Configuring the key server and the KMIP server certificate (in these examples, 1340006-ca.pem), retrieved in step 4.g of Configuring the key server. (In these examples, key_server_9.11.110.162.pem)

   Ensure that the file used in the CLI command exists on the host where the XCLI Utility is installed. If the files downloaded from SafeNet KeySecure are on a different Linux host, the files must be copied to the host where the utility is installed. If needed, a utility, such as SCP (secure copy), may be used to copy the file.

   Create a chained server certificate PEM file by combining the local and server certificates. Before creating the chained certificate, the certificates can be viewed by using the **cat <certificatename>.pem** command.

   Create the chained certificate using the following command:
   Important: The chained certificate starts with the leaf (for example, the server certificate) and ends with the root certificate (the local CA certificate).
   Verify that the chain is setup in the correct order in order to avoid configuration issues.

   For example:

```
cat key_server_9.11.110.162.pem 1340006-ca.pem > 1340006-chained-key_server_9.11.110.162.pem
```

7. Configure the KeySecure key server on the storage system.
   On the host that will be used to run CLI commands, make sure that the storage system chained server certificate (for example, 1340006-chained-key_server_9.11.110.162.pem) created in step 6 of Configuring the storage system exists.

   a. Perform the following CLI commands from the storage system:
      Note:
      - Use a Linux machine to perform this step as the storage system expects the certificate data to be streamed.
      - The default port is 5696. This can be changed if necessary.

      i. **encrypt_keyserver_define** with the following parameters defined (see Defining a key server):
         - name: `globalkey`
         - ipv4: the Thales key server
         - master: `yes`
         - keyserver_type: `KeySecure`
         - certificate: the chained server PEM certificate

      For example:

      ```
      encrypt_keyserver_define name=globalkey ipv4=(THALES KeyServer IP) keyserver_type=KeySecure master=yes
      certificate="$(tr "\n" "*" </root/IBM_Storage_Extended_CLI/1340006-chained-key_server_9.11.110.162.pem)"
      ```

      Note: When using the **encrypt_keyserver_define** command, be sure to use all existing signed storage array certificates..

      ii. Check the key server, using the **encrypt_keyserver_list** command (see Displaying key server status).
      Note: At this point, when encryption on the storage system has still not be enabled, it is okay to see *NOKEY* in the App/Key Status column. This means that there is still no key encrypting the data.

Creating recovery keys

Best practice recommends that recovery keys be configured. When recovery keys are configured, be sure that the command below includes **recovery_keys=yes**.

.For more information regarding recovery keys and how to generate them, see the "Recovery key use and maintenance" section in the *Data-at-rest Encryption for the IBM Spectrum Accelerate Family* Redpaper (REDP-5402).

In order to create recovery keys, at least two Security Admin users must be defined on the storage system.

1. Define at least two Security Admin users on the storage system.
   a. Define each of the Security Admin users, using the **user_define** command (see Defining a new user). For example:

   ```
   user_define user=secadmin1 category=securityadmin password=<password> password_verify=<password>
   user_define user=secadmin2 category=securityadmin password=<password> password_verify=<password>
   ```

   b. Verify the creation of both Security Admins, using the **user_list** command (seeListing users).
2. Create the recovery keys.
   Create two recovery keys, one per Security Admin user.
   a. Define the Security Admin users who are to be receiving the recovery key, using the **encrypt_recovery_key_generate** command (see Generating recovery keys).
   For example:

   ```
   encrypt_recovery_key_generate users=secadmin1,secadmin2  min_req=2
   ```

   b. Log in with each Security Admin credential and generate the recovery key, using the **encrypt_recovery_key_get** command (see Retrieving the security administrator's recovery key).
   The following is an example of a full recovery key:

   ```
   62807CB1902AM074EDLA4EV8F0C574E40A1564F55570CDEEBED37BC3876789
   ```

3. Each Security Admin must verify their recovery key, using the **encrypt_recovery_key_verify** command (see Recovering key verification).

Activating the encryption

1. Update the key server, using the **encrypt_enable recovery_keys=yes -y** command (see Enabling encryption).
2. Run the **state_list** command to check that encryption is *Enabled* (see Listing the operational state). If the status is *Enabling*, run the command again, until status is *Enabled*.
3. Check the key server, using the **encrypt_keyserver_list** command (see Displaying key server status).
   The App/Key Status column should show an *Active* status.

Verifying the status

After the encryption keys have been successfully enabled, the flash enclosures and vault devices start changing from a *Ready* to an *Enrolled* state.

- Verify the encryption state of the flash enclosures and vault devices, by running the following CLI commands as a Security Administrator:
  Note: There are corresponding events visible when using the **event_list** command (see Listing events).
  - **flash_enclosure_list -S** (see Listing flash enclosure status)
  - **vault_device_list -S** (see Listing vault devices in the system)

Migrating from global key to Client Certificate Authentication

Migrate from global key to Client Certificate Authentication.

All procedures that are detailed here take place on the Thales SafeNet KeySecure key server, unless otherwise stated.

Creating a Local User on SafeNet KeySecure

If a user has already been created to be used in the Client Certificate Authentication of other devices, skip this section, and proceed to Editing the AES encryption key properties.

If a unique user for each device is to be used, this step is mandatory for each and for each FlashSystem A9000R storage array.

1. Define a user, from Security > Local Authentication > Local Users & Groups > User & Group Configuration.
   a. Click Add.
   b. Define a username.
      The username must be the same value as the field selected to be the username in the client certificate. Use this same value in Editing the AES encryption key properties
      Note: The Username must begin with a letter.
   c. Enter a password.
   d. Select User Administration Permission and Change Password Permission.
   e. Click Save.

Editing the KMIP key server authentication settings

1. Navigate to KeySecure > Device tab > Key Server > KMIP.
2. Click Properties.
   The Cryptographic Key Server Configuration window opens
3. On the Authentication Settings form, click Edit.
4. Edit the following settings:
   - Change Client Certificate Authentication from *Not Used* to *used for SSL session only*.
   - Select the correct Trusted CA List Profile from the drop-down list.
   - Select the *CN (Common Name)* from the Username Field in Client Certificate drop-down list.

   Click Save.

Editing the AES encryption key properties

1. From Security > Keys > Key List, click Properties.
2. From the Key Properties form, click Edit.
3. Set the Owner Username to the username created in step 1 of Creating a Local User on SafeNet KeySecure, which matches the Username (CN) field.
   This limits access to this specific key to only this specific user.
4. Click Save.

Configuring High Security settings

In order to improve security, disable the use of global keys.

1. On the High Security Configuration from Security > Advanced Security > High Security, scroll to the High Security Settings section and click Edit.
2. Select Disable Creation and Use of Global Keys.
3. Click Save.

Verifying the migration status on the storage system

1. Navigate to the Services Configuration pane from Device > Services > .
2. Select the NAE Server service and click Restart.
3. Click Confirm on the Confirmation Required approval.
4. Check the logs.
5. Check the KeySecure communication status on your storage system, using the following CLI commands:
   a. **encrypt_keyserver_check_status** (see Checking key server status).
   b. **encrypt_keyserver_list** (see Displaying key server status).
      The the App/Key Status status should show the *Active* status.
6. Check the FlashSystem A9000R event log for the following:

```
2020-04-06 20:36:52  349 ENCRYPT_KEYSERVER_ADDED  A key server named 'globalkey' was added.
2020-04-06 20:44:21  351 ENCRYPT_ENABLE_STARTED   Starting encryption activation. This process can take several
minutes to complete.
2020-04-06 20:46:03  360 ENCRYPT_ENABLE_COMPLETED  Encryption is in effect.
2020-04-06 21:03:25  385 ENCRYPT_KEYSERVER_CHECK_STATUS_STARTED  Start checking connectivity status of the
keyservers currently defined in the system. This process can take several minutes to complete.
2020-04-06 21:03:25  388 ENCRYPT_KEYSERVER_CHECK_STATUS_COMPLETED  Completed checking connectivity status of the
keyservers currently defined in the system.
```

7. Check the activity log on the SafeNet KeySecure Management Console.

Thales SafeNet KeySecure 8.11.1 client certificate authentication for key management

Configure an external encryption key management scheme with Thales SafeNet KeySecure 8.11.1 Client Certificate Authentication on your storage system.

The information in this section:

- Does not include installation of the key server.
- Does not cover all available configuration options. It contains the basic configuration to allow encryption between a FlashSystem A9000R storage system and the SafeNet KeySecure key server.
- Uses the storage system management port example IP address `9.11.80.61` and the SafeNet KeySecure server address `9.11.110.162`. Replace these addresses with the correct IP addresses of your FlashSystem A9000R storage system and of your key server.
- Uses the FlashSystem A9000R machine serial number as a prefix when naming PEM files. For example, if the file is saved as 1340006-ca.pem, *134006* indicates the machine serial number of the storage array. To get the serial number of a FlashSystem A9000R storage array, use the **config_get** CLI command (see Displaying the values of configuration parameters).

Known issues

When using SafeNet KeySecure key server with FlashSystem A9000R, and disabling encryption, it may fail with `ENCRYPT_UNABLE_TO_UPDATE_KEY_DURING_DEACTIVATE_ON_KEYSERVER`. This happens if the system was already encrypt enabled, then disabled and enabled again. This can also happen if the key was in Pre-Active state on the key server when encrypt enabling occurred and if this state was not changed later to 'Active'. In order to prevent this issue from occurring, before encryption, enable the key on the key server side. The key on the key server cannot be deactivated or re-activated if it was already deactivated once. If it is deactivated, a new key will be needed on the key server side.

- **Enabling encryption on the storage system**
  Enable encryption on the storage system side with the SafeNet KeySecure.

Enabling encryption on the storage system

Enable encryption on the storage system side with the SafeNet KeySecure.

Preconfigure the storage system by using the following steps:

1. Verify that the storage system is encryption-ready by running the **state_list** command (see Listing the operational state).
   If encryption is not supported, contact IBM Support.
2. Check if any certificates are defined on the system, using the **pki_list** command (see Listing PKI items).
   Note: Having certificates that are predefined on the system does not cause any issues with the configuration.

Configuring the key server

All procedures that are detailed here take place on the Thales SafeNet KeySecure key server, unless otherwise stated.

1. Configure the High Security settings.
   Navigate to Security > High Security and verify that Disable Creation and Use of Global Keys is selected.
2. Create a local Certificate Authority (CA).
   Note:
   - This step might not be necessary if a CA is already in the Local Certificate Authority List.
   - This PEM text example and its decoded X.509 are presented for illustrative purposes only and do not completely fit the values shown in the associated screen captures. For example, there may be inconsistencies in the Common Name (CN) and the Certificate Name field values.

   a. Navigate to Security > Local CAs.
   b. From the Local Certificate Authorities section, populate all of the blank fields.
      Important: Use a strong enough key length and algorithm, such as RSA-2048, according to your company's security policies.
   c. Make sure that Self-signed Root CA is selected.
   d. Click Create.
      The Local Certificate Authority List appears.
   e. Select the CA certificate and click Download.
      This file will be used in later steps to create the client and server chained certificates. (See steps 5 and 6 of Configuring the storage system.)
   f. Save the root (local) CA certificate in PEM format. For example, 1340006-ca.pem.
      The PEM format encodes an X.509 certificate.
   g. Check that the root certificate can sign other certificates. Use the following command:

      ```
      # openssl x509 -in 1340006-ca.pem -inform pem -text -noout
      ```

      Note: The following are indications of a root certificate which can sign other certificates:
      - The Issuer and the Subject are the same.
      - There is an X.509 extension indicating it is a CA (**`CA:TRUE`**).
3. Add the local CA to the KeySecure's trusted CA list.

a. Navigate to Security > Trusted CA Lists.
b. From the Trusted Certificate Authority List Profiles form, select an existing profile for KMIP connections or create a new one.
   - Select the existing Default profile and click Properties.
   - Add a specialized profile for KMIP connections.
     Create the new profile by clicking Add. Enter a name for the profile and click Save. Select the profile and click Properties.

c. From the Trusted Certificate Authority List, click Edit.
   The Trusted CAs and Available CAs lists are displayed.

d. Select the local certificate authority and move it to the Trusted CAs list, by clicking Add.
e. Click Save.

4. Create an SSL Certificate for the KMIP key server and sign it with the local CA.
   Note: This step might not be necessary if the customer environment already has an SSL Server Certificate configured for the KeySecure key server.
   a. Create an SSL Certificate for the KMIP key server and sign it with the local CA.
      i. Navigate to Security > SSL Certificates.
      ii. From the Create Certificate Request pane, populate all of the blank fields.
      iii. Click Create Certificate Request.
   b. Verify that the certificate was added from the Certificate List. Select the SSL certificate and click Properties.
   c. Copy the entire text of the SSL certificate, including the `-----BEGIN CERTIFICATE REQUEST-----` and `-----END CERTIFICATE REQUEST-----` text.
   d. Go to Security > Local CAs and click Sign Request.
      The Sign Certificate Request form opens.

   e. Paste the SSL certificate text, copied in step 4.c into the Certificate Request box.
   f. From the Certificate Purpose options, select Server and click Sign Request.
   g. From the CA Certificate Information form, copy the entire certificate text, including the `-----BEGIN CERTIFICATE REQUEST-----` and `-----END CERTIFICATE REQUEST-----` text.
   h. Go to Security > SSL Certificates and click Properties.
      The Certificate Request Information page opens.

   i. At the bottom of the page, click Install Certificate.
   j. Paste the certificate text copied in step 4.g and click Save.
   k. From Security > SSL Certificates > Certificate and CA Configuration verify that the Certificate Status in the Certificate List is *Active*.
   l. Click Properties.
   m. Save the server certificate in PEM format by clicking Download from the Certificate Information form.
   n. Save the server certificate in PEM format. For example, key_server_9.11.110.162.pem.
      Note: Up to four key servers can be configured at the same time on each FlashSystem A9000R storage array and each one must be configured with its own certificate that is signed by the same CA certificate. Therefore, be sure to identify each server PEM file with the IP address or FQDN of the specific key server for easy identification. When using the **encrypt_keyserver_define** command, be sure to use all existing signed storage array certificates. The PEM file is used in 6 of Configuring the storage system.

      The PEM format encodes an X.509 certificate.

   o. Check that the root certificate can sign other certificates.
      Use the following command:

      ```
      # openssl x509 -in key_server_9.11.110.162.pem -inform pem -text -noout
      ```

      Note: The following are indications of a signed certificate which cannot sign other certificates:
      - The Issuer and the Subject are *not* the same.
      - There is an X.509 extension indicating it is *not* a CA (`CA:FALSE`).

5. Add a KMIP key server.
   a. Navigate to Device > Key Server and click Add.
      Note: Only the KMIP server should remain. This is a strong security recommendation.
   b. Select and enter the following properties within the Cryptographic Key Server Properties form:
      - Protocol: *KMIP*
      - Port: `5696`
      - Select Use SSL.
      - Select the Server Certificate defined in step 4 of Configuring the key server from the drop-down list.
   c. Click Save.
   d. Select the KMIP protocol and click Properties.
      The Cryptographic Key Server Configuration page appears.

   e. From the Authentication Settings from, click Edit.
   f. From KeySecure > Device tab > Key Server > KMIP edit the KMIP key server Authentication Settings.
      - Change Client Certificate Authentication from *Not Used* to *used for SSL session only*.
      - Select Default from the Trusted CA List Profile drop-down list.
      - Select CN (Common Name) from the Username Field in Client Certificate drop-down list.

      Click Save.

6. Create an AES encryption key.
   a. Navigate to Security > Keys > Key List.
   b. Click Create Key to add an AES encryption key.
   c. Set the Create Key fields.
      - Fill in the Key Name. The key name is the machine model and serial number of your storage system. For example, `425-1340006`, where *425* is the machine model and *1340006* is the serial number.
      - Fill in the Owner Username field, using the same one used in step 7.c and 1 of Configuring the storage system blank.
      - From the Algorithm drop-down list, select AES-256.
      - Select Deletable and Exportable.

- Select Immediately in the Activation Date field.

Click Create.

7. Create a local user on the KeySecure key server
   Note: If a user has already been created to be used in the Client Certificate Authentication of other devices, this step may be skipped. If you want to use a unique user for each device, this step is mandatory for each and every device. (For example, for each FlashSystem A9000R storage array).
   a. Navigate to the Local Users section of Security > Local Authentication > User & Group Configuration.
   b. Click Add.
   c. Enter a Username.
      The Username must be the same value as the field selected to be the username in the client certificate. Use this same value when creating the Certificate Signing Request (CSR) of for the KMIP client (the FlashSystem A9000R storage array). (See step 1 of Configuring the storage system).
   d. Enter a Password
   e. Click Save.

Configuring the storage system

1. Create the Certificate Signing Request (CSR) for the KMIP client certificate.
   a. Create the Certificate Signing Request (CSR) on the storage system side.
      This generates a certificate, which is later used to sign with local CA in the SafeNet KeySecure server (see step 2 of Configuring the storage system).
      Note: The CN part should be the same as the SafeNet local username created previously in step 7.c of Configuring the key server.
      Define the following parameters in the **`pki_generate_private_key_and_csr`** command (see Generating a private key and CSR):
      - Name is the certificate's symbolic name.
      - Subject is the subject name for the generated certificate request.

      The argument must be formatted as **`/type0=value0/type1=value1/type2=...`** .
      Include the email of the person who needs to be notified of any certificate or security issues.

      For example,

      ```
      pki_generate_private_key_and_csr name=6003306_client subject="/C=US/ST=Arizona/L=Tucson/O=IBM/OU=IBM-lab/CN=user/emailAddress=name@company.com"
      CSR
      ```

      where *csrforks* is the name and *user* is the CN (Common Name) username, with the storage array located in Tuscon, Arizona.
      The storage system certificate is displayed.

   b. Copy the entire text of the storage system SSL Certificate Signing Request (CSR), including the **`'-----BEGIN CERTIFICATE REQUEST-----' and '-----END CERTIFICATE REQUEST-----'`**.
2. Sign and install the storage system Certificate Signing Request (CSR).
   Use a KeySecure key secure local CA to sign the Client Certificate Request.
   a. Navigate to Security > Local CAs > Local Certificate Authority List.
   b. Select a local CA created in step 2 of Configuring the key server and click Sign Request.
   c. On the Sign Certificate Request form:
      i. Paste the certificate request copied in step 1.b into the Certificate Request field.
      ii. From the Certificate Purpose options, select Client.
      iii. Specify a Certificate Duration.

   Click Sign Request.
3. Download the signed storage system client certificate PEM file.
   a. Click Download to download the certificate.
   b. Save the signed client certificate in PEM format. For example: 1340006-client.pem
      The PEM format encodes an X.509 certificate.

4. Check that the client certificate cannot sign other certificates (that it is a leaf certificate).
   Use the following command:

   ```
   # openssl x509 -in 1340006-client.pem -inform pem -text -noout
   ```

   Note: The following are indications of a root certificate which can sign other certificates:
   - The Issuer and the Subject are the same.
   - There is an X.509 extension indicating it is a CA (**`CA:TRUE`**).
5. Create a chained client certificate PEM file.
   These steps must be done on the Linux host where the two files were downloaded from the key server.

   Use the local root CA PEM file, created and downloaded in step 2 of Configuring the key server (in these examples, 1340006-ca.pem) and the signed storage client certificate PEM file, downloaded in step 3 of Configuring the storage system (in these examples, 1340006-client.pem).

   Ensure that the file used in the CLI command exists on the host where the XCLI Utility is installed. If the files downloaded from SafeNet KeySecure are on a different Linux host, the files must be copied to the host where the utility is installed. If needed, a utility, such as SCP (secure copy), may be used to copy the file.

   a. Create a chained client certificate PEM file by combining the local and client certificates.
      Before creating the chained certificate, the certificates can be viewed by using the **cat <certificatename>.pem** command.

      Create the chained certificate using the following command:
      Important: The chained certificate starts with the leaf (for example, the storage system client certificate) and ends with the root certificate (the local CA certificate).
      Verify that the chain is setup in the correct order in order to avoid configuration issues.

      For example:

      ```
      cat 1340006-client.pem 1340006-ca.pem >1340006-chained-client.pem
      ```

b. Install the chained client certificate on the storage system and update the certificate's related services.
From the storage system XCLI Utility:
i. Import the signed PEM certificate, using the **pki_set_pem** command (see Importing a signed certificate).

```
pki_set_pem services=none certificate="$(tr "\n" "*" </root/IBM_Storage_Extended_CLI/1340006-chained-client)"
```

ii. List the PKI items, using the **pki_list** command.
iii. Remove the KMIP service from the services list of the chained client certificate, using the following **pki_update** CLI command (see Updating a PKI certificate or services).
The expected output is `Command completed successfully.`

iv. Add the KMIP service to the csrforks, using the **pki_update** command.
v. Verify that all PKI items are properly updated, using the **pki_list** command.

6. Create a chained server certificate PEM file.
These steps must be done on the Linux host where the two files were downloaded from the key server.

Use the local root CA PEM file, created and downloaded in step 2 of Configuring the key server and the KMIP server certificate (in these examples, 1340006-ca.pem), retrieved in step 4.g of Configuring the key server. (In these examples, key_server_9.11.110.162.pem)

Ensure that the file used in the CLI command exists on the host where the XCLI Utility is installed. If the files downloaded from SafeNet KeySecure are on a different Linux host, the files must be copied to the host where the utility is installed. If needed, a utility, such as SCP (secure copy), may be used to copy the file.

Create a chained server certificate PEM file by combining the local and server certificates. Before creating the chained certificate, the certificates can be viewed by using the **cat <certificatename>.pem** command.

Create the chained certificate using the following command:
Important: The chained certificate starts with the leaf (for example, the server certificate) and ends with the root certificate (the local CA certificate). Verify that the chain is setup in the correct order in order to avoid configuration issues.

```
cat key_server_9.11.110.162.pem 1340006-ca.pem > 1340006-chained_key_server_9.11.110.162.pem
```

7. Configure the KeySecure key server on the storage system.
On the host that will be used to run CLI commands, make sure that the storage system chained server certificate (for example, 1340006-chained-key_server_9.11.110.162.pem) created in step 6 of Configuring the storage system exists.

a. Perform the following CLI commands from the storage system:
Note:
- Use a Linux machine to perform this step as the storage system expects the certificate data to be streamed.
- The default port is 5696. This can be changed if necessary.

i. **encrypt_keyserver_define** with the following parameters defined (see Defining a key server):
- name: `KS`
- ipv4: the Thales key server
- master: `yes`
- keyserver_type: `KeySecure`
- certificate: the chained server PEM certificate

For example:

```
encrypt_keyserver_define name=KS ipv4=(THALES KeyServer IP) keyserver_type=KeySecure master=yes
certificate="$(tr "\n" "*" </root/IBM_Storage_Extended_CLI/1340006-chained-key_server_9.11.110.162.pem)"
```

Note: When using the **encrypt_keyserver_define** command, be sure to use all existing signed storage array certificates.
ii. If an update is needed, perform the **encrypt_keyserver_update** command (see Displaying key server status). For example:

```
encrypt_keyserver_update name=KS certificate="$(tr "\n" "*" </root/1340006-
chained_key_server_9.11.110.162.pem
```

iii. Check the key server, using the **encrypt_keyserver_list** command (see Displaying key server status).
Note: At this point, when encryption on the storage system has still not be enabled, it is okay to see *NOKEY* in the App/Key Status column. This means that there is still no key encrypting the data.

Creating recovery keys

Best practice recommends that recovery keys be configured. When recovery keys are configured, be sure that the command below includes **recovery_keys=yes**.

.For more information regarding recovery keys and how to generate them, see the "Recovery key use and maintenance" section in the *Data-at-rest Encryption for the IBM Spectrum Accelerate Family* Redpaper (REDP-5402).

In order to create recovery keys, at least two Security Admin users must be defined on the storage system.

1. Define at least two Security Admin users on the storage system.
a. Define each of the Security Admin users, using the **user_define** command (see Defining a new user). For example:

```
user_define user=secadmin1 category=securityadmin password=<password> password_verify=<password>
user_define user=secadmin2 category=securityadmin password=<password> password_verify=<password>
```

b. Verify the creation of both Security Admins, using the **user_list** command (seeListing users).
2. Create the recovery keys.
Create two recovery keys, one per Security Admin user.
a. Define the Security Admin users who are to be receiving the recovery key, using the **encrypt_recovery_key_generate** command (see Generating recovery keys).

For example:

```
encrypt_recovery_key_generate users=secadmin1,secadmin2  min_req=2
```

b. Log in with each Security Admin credential and generate the recovery key, using the **encrypt_recovery_key_get** command (see Retrieving the security administrator's recovery key).
The following is an example of a full recovery key:

```
62807CB1902AM074EDLA4EV8F0C574E40A1564F55570CDEEBED37BC3876789
```

3. Each Security Admin must verify their recovery key, using the **encrypt_recovery_key_verify** command (see Recovering key verification).

Activating the encryption

1. Update the key server, using the **encrypt_enable recovery_keys=yes -y** command (see Enabling encryption).
2. Run the **state_list** command to check that encryption is *Enabled* (see Listing the operational state). If the status is *Enabling*, run the command again, until status is *Enabled*.

Verifying the status

After the encryption keys have been successfully enabled, the flash enclosures and vault devices start changing from a *Ready* to an *Enrolled* state.

- Verify the encryption state of the flash enclosures and vault devices, by running the following CLI commands as a Security Administrator:
  Note: There are corresponding events visible when using the **event_list** command (see Listing events).
  - **flash_enclosure_list -S** (see Listing flash enclosure status)
  - **vault_device_list -S** (see Listing vault devices in the system)

Reference

Refer to the following sections for information about how to use the IBM FlashSystem® A9000R command-line interface (CLI), IBM® Hyper-Scale Manager Representational State Transfer (REST) Application Programming Interface (API), and Open API.

- **Command-line reference (12.3.2.e)**
  The following information describes how to use the IBM FlashSystem A9000R 12.3.2.e command-line interface (CLI).
- **REST API reference**
  The IBM Hyper-Scale Manager REST Application Programming Interface (API) is documented within the IBM Hyper-Scale Manager documentation set.
- **Open API reference**
  The following information describes how to use the Open Application Programming Interface (API) for IBM FlashSystem A9000 and A9000R systems.

Command-line reference (12.3.2.e)

The following information describes how to use the IBM FlashSystem® A9000R 12.3.2.e command-line interface (CLI).

Note: The information for version 12.3.2.e is also applicable to versions 12.3.2.a, 12.3.2.b, 12.3.2.c, and 12.3.2.d.

- **Overview of the XCLI utility**
  This section describes how to download, install, and start the IBM XCLI utility. It also explains how to log off the XCLI.
- **Host and cluster management commands**
  This section describes the command-line interface (CLI) for host and cluster management.
- **Volume management commands**
  This section describes the command-line interface (CLI) for snapshot management.
- **Volume snapshot management commands**
  This section describes the command-line interface (CLI) for snapshot management.
- **Consistency group management commands**
  This section describes the command-line interface (CLI) for consistency group management.
- **Snapshot set management commands**
  This section describes the command-line interface (CLI) for snapshot set management.
- **Storage pool management commands**
  This section describes the command-line interface (CLI) for storage pool management.
- **System management commands**
  This section describes the command-line interface (CLI) for system management.
- **Remote target connectivity commands**
  This section describes the command-line interface (CLI) for defining remote target connectivity.
- **Remote mirroring commands**
  This section describes the command-line interface (CLI) for remote mirroring.

- **HyperSwap commands**
  This chapter describes the command line interface (CLI) for HyperSwap.
- **Data migration commands**
  This section describes the command-line interface (CLI) for data migration.
- **IBM Hyper-Scale Mobility commands**
  This section describes the command-line interface (CLI) for IBM Hyper-Scale Mobility.
- **Event handling commands**
  This section describes the command-line interface (CLI) for event handling, including listing events, filtering and sending notifications.
- **IP configuration commands**
  This section describes the command-line interface (CLI) for IP configuration.
- **PKI configuration commands**
  This section describes the command-line interface (CLI) for PKI configuration.
- **InfiniBand commands**
  This section describes the command-line interface (CLI) for InfiniBand fabric management.
- **Access control commands**
  This section describes the command-line interface (CLI) for user access control.
- **Fibre Channel and iSCSI configuration and status commands**
  This section describes the command-line interface (CLI) for Fibre Channel port configuration.
- **Flash enclosure maintenance commands**
  This section describes the command-line interface (CLI) for maintaining flash enclosures.
- **Hardware maintenance commands**
  This section describes the command-line interface (CLI) for maintaining hardware components.
- **Statistics commands**
  This section describes the command-line interface (CLI) for getting system statistics.
- **Metadata commands**
  This section describes the command-line interface (CLI) for handling metadata.
- **Encryption enablement and support commands**
  This section describes the command-line interface (CLI) for encryption configuration.
- **Multi-site high availability and disaster recovery commands**
  This section describes the command-line interface (CLI) for Multi-site high availability and disaster recovery.
- **Security configuration commands**
  This chapter describes the command line interface (CLI) for security configuration.
- **Events**
- **Return codes**
  This section contains descriptions of CLI return codes.

Overview of the command-line interface (CLI)

The command-line interface (CLI) provides a mechanism for issuing commands to manage and maintain the storage system. CLI commands are entered on the IBM XCLI utility.

This section explains how to install and start the XCLI utility. It also provides information about interactive and basic modes for running commands in the utility and an overview of the CLI command structure and parameters.

The following topics are covered:

- **Overview of the XCLI utility**
  This section describes how to download, install, and start the IBM XCLI utility. It also explains how to log off the XCLI.
- **Using the CLI**
  This section provides information about interactive and basic modes for running CLI commands and an overview of the CLI command structure and parameters.

Overview of the XCLI utility

This section describes how to download, install, and start the IBM XCLI utility. It also explains how to log off the XCLI.

The following topics are covered:

- **Installing and starting the XCLI**
  This information describes how to download and install the IBM XCLI utility.

- **Exiting an interactive XCLI session**
  When running XCLI in interactive mode, you can exit the interactive session by either closing the session window or running the **exit** command.

Installing and starting the XCLI

This information describes how to download and install the IBM XCLI utility.

About this task

Note: For the installation requirements and a list of available packages, see the *XCLI Utility Release Notes*.

Procedure

Perform these steps to download and install the XCLI:

1. Download the IBM XCLI installation package from the IBM Fix Central website.
2. Perform one of the following procedures for your operating system.
   Note: See the IBM XCLI Utility release notes for your version to see which operating systems are supported. (See Release note PDF publications).
   - **Windows:** Double-click the installation file, and follow the instructions on the screen.
   - **Linux®, AIX®, HP-UX, Solaris:** Run the installation file using the following command:

     ```
     ./<filename>
     ```

     Then set the program path.
3. Start the XCLI depending on the hosting operating system and operational mode.

- **Starting the XCLI on a Windows system**
  You can start the XCLI on a Windows system in either interactive or basic mode.
- **Starting the XCLI on a Linux system**
  You can start the XCLI on a Linux system in either interactive or basic mode.

Starting the XCLI on a Windows system

You can start the XCLI on a Windows system in either interactive or basic mode.

Interactive mode

About this task

To run commands in interactive mode, perform the following steps:

Procedure

1. Click Start > Programs > IBM XIV > XCLI to open an XCLI session window.
2. Follow the instructions on the screen and type the following information:
   a. Storage system IP address or DNS
   b. User name
   c. Password

   Figure 1. XCLI terminal

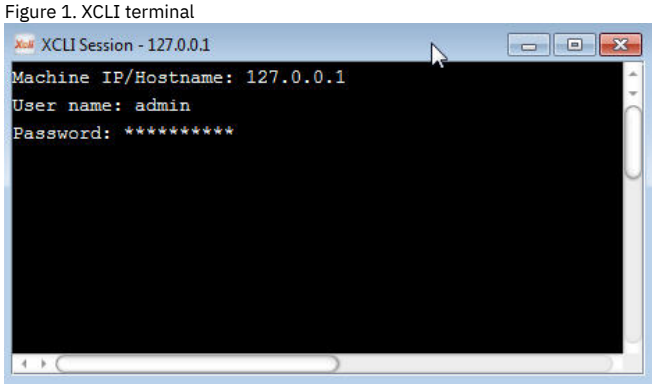

3. You are now connected to the specified storage system. The XCLI prompt appears in the session window. The window title includes the name of the storage system to which the XCLI is connected.
4. Run any CLI command from this prompt.
   Note: The font size in the XCLI utility terminal is customizable. To increase the font size, press CTRL+. To decrease the font size, press CTRL-.

Basic mode

About this task

To run commands in basic mode, perform the following steps:

Procedure

1. Open a Windows command session.
2. Type `cd c:\program files\IBM\Storage\XIV`.
3. Run any CLI command, including the XCLI identification parameters, for example:

```
xcli -u user -p ******** -m 127.0.0.1 vol_list
```

Starting the XCLI on a Linux system

You can start the XCLI on a Linux® system in either interactive or basic mode.

Interactive mode

About this task

To run commands in interactive mode, perform the following steps:

Procedure

1. Type **`xcli -w`** to open a new session.
2. Follow the instructions on the screen and type the following information:
   a. User name
   b. Password
   c. Storage system IP address or DNS
3. You are now connected to the specified storage system. The XCLI prompt appears in the session window. You can run any CLI command from this prompt.

Basic mode

About this task

To run commands in basic mode, enter the command including the XCLI identification parameters, for example:

```
xcli -u user -p ******** -m 127.0.0.1 vol_list
```

Exiting an interactive XCLI session

When running XCLI in interactive mode, you can exit the interactive session by either closing the session window or running the **exit** command.

To terminate an interactive XCLI session automatically after the predefined period, set the session_timeout property in the xiv-general.properties file, that is stored in the properties directory. The value of session_timeout is defined in minutes.

The location of the **`properties`** directory varies among operating systems as follows:

Windows
`%APPDATA%\XIV\GUI12\properties`
Non Windows
`$USERDIR/.xiv/GUI12/properties`

Using the CLI

This section provides information about interactive and basic modes for running CLI commands and an overview of the CLI command structure and parameters.

The following topics are covered:

- **Interactive and basic modes**
  You can use the CLI in two modes: interactive and basic.
- **Understanding the command syntax**
  This information describes the general syntax for a CLI command in basic mode.
- **Using identification parameters**
  This information describes the parameters used to set the user, password, and storage system.
- **Display options**
  This information describes the formats that you can choose to display the command output.

- **Running commands in batch mode**
  CLI commands can be grouped together and run in a batch. For example, you can use batch mode to run an identical set of commands on multiple storage systems.
- **Displaying CLI help and the XCLI utility version**
  This information describes how to display help for the CLI command and the version of the XCLI utility.

Interactive and basic modes

You can use the CLI in two modes: interactive and basic.

The differences between these two modes are as follows:

- Basic mode requires you to log in to the storage system each time you issue a command, but the interactive mode requires you to log in only once.
- In basic mode, you must enter the entire command syntax. In interactive mode, you can enter a shorter syntax.
- Interactive mode offers several command and argument completion features.

The following example shows the command syntax for each of these modes:

Basic

```
xcli -u user -p password -m 127.0.0.1 vol_list
```

Interactive

```
vol_list
```

Note: Each of the batch-mode parameters in these examples is explained later in this chapter.

- **Interactive mode features**
  Running CLI commands in interactive mode provides command and argument completion, along with possible values to these arguments.

Interactive mode features

Running CLI commands in interactive mode provides command and argument completion, along with possible values to these arguments.

About this task

The CLI offers several ways to interactively complete command names.

Procedure

1. To complete the name of a command, type one or more characters and press Tab.
   In the following example, the CLI adds a character to the name of a command that starts with vol. The first click on Tab adds a character.

   ```
   >>vol
   >>vol_
   ```

2. Next, to list commands, type one or more characters and press Tab twice (Tab-Tab).

   ```
   >>vol_
   vol_by_id          vol_copy           vol_create         vol_delete
   vol_format         vol_list           vol_lock            vol_mapping_list
   vol_move           vol_rename         vol_resize         vol_unlock
   ```

Example

- This example lists all commands that start with the letter v:

  ```
  >> v
  version_get        vol_by_id          vol_copy           vol_create
  vol_delete         vol_format         vol_list           vol_lock
  vol_mapping_list   vol_move           vol_rename         vol_resize
  vol_unlock         vpd_config_get     vpd_config_set
  ```

- This example lists all commands that start with the letters cg_:

  ```
  >> cg_
  cg_add_vol         cg_create          cg_delete          cg_list
  cg_move            cg_remove_vol      cg_rename          cg_snapshots_create
  ```

Command Argument Completion

The CLI can provide list arguments and argument values to help you complete a command.

About this task

To list arguments for a specific command, type the command name and press Tab. To list values for a specific argument, type the command name and argument, optionally specify one or more characters for the value, and then press Tab.

Procedure

1. Listing arguments for a command.
   This example lists arguments for the vol_create command:

   ```
   >> vol_create
   vol=   size=   size_block=   pool=
   ```

2. Listing values for a specific argument.
   This example lists valid values for the pool argument that is required for the pool_create command:

   ```
   >> pool_create pool=
   8058  2nd Pool  8059  pool1
   ```

3. Listing a subset of values for a specific argument.
   This example lists valid values that start with v for the vol argument that is required for the vol_list command:

   ```
   >> vol_list vol=v
   vol1  vol2
   ```

Understanding the command syntax

This information describes the general syntax for a CLI command in basic mode.

When in basic mode, the CLI uses the following general syntax:

```
xcli < <[ --file=FILE ] -c CONFIG|-m IP1 [-m IP2 [-m IP3]] >
-L|-a <config> -m IP1 [-m IP2 [ -m IP3]]|-d IP1 [-m IP2 [ -m IP3]] >|
[ -r ]
[ <-l | --list> | <-s | --csv> | <-x | --xml> ]
[ <-u | --user> user ]
[ <-p | --password> password ]
[-t [--fields=field1,field2,...]]
[command]
```

- **Command syntax quick reference**
  This information describes the command parameters and options that are available in the CLI basic mode.
- **Syntax example**
  The CLI command syntax specifies the command to be run, along with its applicable parameters and their values.

Command syntax quick reference

This information describes the command parameters and options that are available in the CLI basic mode.

Use the following table as a quick reference to the various parameters and options.

Options	Values
-f	Specifies the name of a configuration file that lists the storage system
-c	Specifies the storage system on which the command is to be run
-m	Specifies the IP address of the storage system on which the command runs
-L	Lists the storage systems, as read from the configuration file
-a	Specifies the name of the storage system
-d	Removes a storage system from the configuration file
-r	Specifies the name of a batch file that runs CLI commands
-l	Displays the command output in user-readable format
-s	Displays the command output in CSV format
-x	Displays the command output in XML format
-u	Specifies the user
-p	Specifies the password
-t	Manages the fields of the command output
-h	Displays command help
-y	Suppresses the **Are you sure?** prompt
-v	Displays the version of the XCLI on the screen
command	Runs the specified command

Syntax example

The CLI command syntax specifies the command to be run, along with its applicable parameters and their values.

In the following example, the parameters to the left of the **vol_list** command specify the storage system to which the command is being directed, and also specify the required user and password for this storage system:

```
xcli -u admin -p ********** -m 127.0.0.1 vol_list
```

Identification parameters

-u
: Specifies the user ID.

-p
: Specifies the password.

Storage system
: The storage system is specified by either its IP address or name of the storage system as listed in the configuration file. See Configuration parameters for more information.

-m
: Specifies the IP address of the storage system to which this command is directed.

-c
: Specifies the name of the storage system to which this command is directed, as it is defined in the configuration file (for example, `my_system`).

Command

vol_list
: Specifies the command to be run. For more information about running commands, see Interactive mode features.

Using identification parameters

This information describes the parameters used to set the user, password, and storage system.

The following topics are covered:

- **Setting user and password parameters**
  The CLI and the storage system provide a password-controlled user ID as a security mechanism for controlling CLI operations.
- **Identifying and configuring a storage system**
  This information describes the parameters used to identify the storage system on which a command is to run, and how to create a configuration file to manage the storage systems that you can use.
- **Certificate management**
  This section describes the way certificates are managed via the XCLI utility.

Setting user and password parameters

The CLI and the storage system provide a password-controlled user ID as a security mechanism for controlling CLI operations.

When running in basic mode, specify the user name and password as follows:

```
xcli -u admin -p ********** -c my_system vol_list
```

In this command:

Identification parameters

-u
: Specifies the user ID.

-p
: Specifies the password.

my_system

-c
: Specifies the name of the storage system to which this command is directed, as it is defined in the configuration file (for example, `my_system`).

Command

vol_list
: Specifies the command to be run.

The password handling mechanism performs as follows:

1. Checking the user:
   - The -u or --user parameter on the command line is checked first and its value is used as the user name.
   - If the -u or --user parameter is not specified, the **XIV_XCLIUSER** environment variable is used as a user name.
2. Checking the password:
   - The -p or --password parameter on the command line is checked first and its value is used as the password.
   - If the -p or --password parameter is not specified, the **XIV_XCLIPASSWORD** environment variable is used as the password.

Note: If you do not specify both the user ID and the password, the command fails.

Identifying and configuring a storage system

This information describes the parameters used to identify the storage system on which a command is to run, and how to create a configuration file to manage the storage systems that you can use.

- **Configuration parameters**
  Most CLI commands are directed to a specific storage system using the IP address. You must provide at least one address and up to three addresses per storage system.
- **Creating a configuration file**
  You can use the configuration file to manage a list of the storage systems that you are working with.

Configuration parameters

Most CLI commands are directed to a specific storage system using the IP address. You must provide at least one address and up to three addresses per storage system.

To provide the storage system IP address, log in to an interactive session or specify the configuration file that stores the storage system IP address or addresses.

Specifying a storage system using its IP address
In the following example, the command is directed to a storage system with an IP address of `127.0.0.1`:

```
xcli -u admin -p ********** -m 127.0.0.1 vol_list
```

In this command:

Identification parameters

-u
Specifies the user ID.
-p
Specifies the password.

Storage system
The storage system is specified by its IP address.

-m
Specifies the IP address of the storage system to which this command is directed.

Command

vol_list
Specifies the command to be run.

Specifying a storage system by using a configuration file
In the following example, the command is directed to a storage system that is listed on the `my_system` configuration file:

```
xcli -u admin -p ********** -c my_system vol_list
```

In this command:

Identification parameters

-u
Specifies the user ID.
-p
Specifies the password.

Storage system
The storage system is specified by its name on the configuration file.

-c
Specifies the name of the storage system to which this command is directed, as it is defined in the configuration file (for example, `my_system`).

Command

vol_list
Specifies the command to be run.

Creating a configuration file

You can use the configuration file to manage a list of the storage systems that you are working with.

Use the following options to add and subtract storage systems from this file and to list them.

Listing the available storage systems
: In the following example, the configuration information is read from a default file location or from the file that is specified with **[-f *file*]**.

```
xcli [-f file] -L
```

Adding a new storage system to the configuration file
: In the following example, IP1...IP3 are added to the configuration file at the default file location. If applicable, the addresses are added to the file that is specified in `[-f file]`. The `<config>` variable represents the configuration name of the storage system that you are adding to the list.

```
xcli [-f file] -a <config> -m IP1 [-m IP2 [ -m IP3]]
```

Removing a storage system from the configuration file
: In the following example, IP1...IP3 are removed from the configuration file. If applicable, the addresses are removed from the file that is specified with `[-f file]`.

```
xcli [-f file] -d IP1 [-m IP2 [ -m IP3]]
```

Location of the configuration file

The configuration file is located in the following directory, depending on the operating system. You do not specify the location of the configuration file when you add or remove storage systems from the configuration.

Windows
: \Application Data\XIV\GUI12\properties

UNIX
: In the home folder under .xiv

Certificate management

This section describes the way certificates are managed via the XCLI utility.

The general format of the certificate commands is:

```
xcli -C <command> [ <p1>=<v1> [<p2>=<v2>]...]
```

The available commands are: list, show, import and remove.

List [type=<type>]

This command lists the trusted certificates (global and private). This command accepts the type of list as a parameter.

Type = `all` (default)
: Lists all trusted certificates.
For example:

```
xcli -C list
```

Private
: Lists all private trusted certificates.
For example:

```
xcli -C list type=private
```

Global
: Lists all global trusted certificates.

Show alias=<alias>

This command displays the certificate details. This command accepts the name of the specific certificate as a parameter. For example:

```
xcli -C show alias=abcd
```

Import pem=<pem_file_path> [alias=<alias>]

This command imports a certificate (in PEM format) into the list of trusted certificates. This command accepts the location of the certificate as a mandatory parameter and the name into which the certificate will be renamed. For example:

```
xcli -C import pem=C:\abc\def\cert.pem
xcli -C import alias=abcd pem=C:\abc\def\cert.pem
```

Remove alias=<alias>

This command removes a certificate from the list. For example:

```
xcli -C remove alias=abcd
```

Display options

This information describes the formats that you can choose to display the command output.

The following topics are covered:

- **Using display options**
  Output from an CLI command can be displayed in a list, comma-separated value (CSV) and XML formats. You can specify only one format. If you do not specify the format, the output defaults to a list.
- **Table display options**
  The list option displays the command output in a user-readable format. When running a command with a list option, you can specify which table columns are displayed on the screen.

Using display options

Output from an CLI command can be displayed in a list, comma-separated value (CSV) and XML formats. You can specify only one format. If you do not specify the format, the output defaults to a list.

The display options are:

-l
: Displays command output in a list (also known as user-readable format).

-s
: Displays command output in CSV format.

-x
: Displays command output in XML format.

Use the display options as follows:

Interactive mode

```
vol_list -s
```

Basic mode

```
xcli -u user -p ********** -m 127.0.0.1 -s vol_list
```

In this command:

Identification parameters

-u
: Specifies the user ID.

-p
: Specifies the password.

Storage system
: The storage system is specified by either its IP address or name of the storage system as listed in the configuration file. See Configuration parameters for more information.

-m
: Specifies the IP address of the storage system to which this command is directed.

-c
: Specifies the name of the storage system to which this command is directed, as it is defined in the configuration file (for example, `my_system`).

Display option

-s
: Displays command output in CSV format.

Command

vol_list
: Specifies the command to be run. For more information about running commands, see Interactive mode features.

Table display options

The list option displays the command output in a user-readable format. When running a command with a list option, you can specify which table columns are displayed on the screen.

Determine the way that the table is displayed as follows:

Interactive mode

```
vol_list -f "size"
```

Single-command mode

```
xcli -u admin -p ********** -m 127.0.0.1 vol_list -f "size"
```

In this command:

Identification parameters

-u
: Specifies the user ID.

-p
: Specifies the password.

Storage system
: The storage system is specified by either its IP address or name of the storage system as listed in the configuration file. See Configuration parameters for more information.

-m
: Specifies the IP address of the storage system to which this command is directed.

-c
: Specifies the name of the storage system to which this command is directed, as it is defined in the configuration file (for example, `my_system`).

Table display option

-f "size"
: Specifies the columns to be displayed. Multiple columns can be specified by a comma-separated list.
In this example, only the Size column is displayed. You can list any combination of the table columns.

Command

vol_list
: Specifies the command to be run. For more information about running commands, see Interactive mode features.

Viewing the available columns

You can view all of the available table's columns by running: **xcli.py help command=<*command_name*> -f fields -z**.

The result provides information about the command, including a list of all of its output fields.

Running commands in batch mode

CLI commands can be grouped together and run in a batch. For example, you can use batch mode to run an identical set of commands on multiple storage systems.

Creating a batch file for the commands

Create a text file and write the commands without the **xcli** prefix or CLI parameters. For example:

```
pool_create pool=pool_00001 hard_size=171 soft_size=171 snapshot_size=65
vol_create vol=vol_00010 size=17 pool=pool_00001
vol_list vol=vol_00010
```

This example contains the following commands:

- The **pool_create** command, along with its arguments. This command creates a storage pool, which is a prerequisite for creating a volume.
- The **vol_create** command, along with its arguments. This command creates a volume in the pool that has just been created.
- The **vol_list** command displays the details of the newly created volume.

Name the script file and save it.

Running a batch file

To run the batch file, you must specify the CLI parameters:

```
xcli -u admin -p ********** -m 127.0.0.1 -r
"C:\Documents and Settings\avia\xcli\xcli_script.txt"
```

In this command:

Identification parameters

-u
: Specifies the user ID.

-p
: Specifies the password.

Storage system
The storage system is specified by either its IP address or name of the storage system as listed in the configuration file. See Configuration parameters for more information.

-m
Specifies the IP address of the storage system to which this command is directed.
-c
Specifies the name of the storage system to which this command is directed, as it is defined in the configuration file (for example, `my_system`).

The batch parameter

-r
Specifies the name of the batch file to run on the storage system.

Viewing the output

The three commands in the previous example create a pool, then create a volume, then display the volume details. The following output is returned from running these three commands in batch mode:

- Confirmation that a pool was created
- Confirmation that a volume was created
- Table with the details of the newly created volume

Failure of batch mode

When one of the commands that run in batch mode fails, the following actions occur:

1. The script exits immediately.
2. No commands after the failing command are run.
3. An error message is displayed identifying the CLI command that failed.

Displaying CLI help and the XCLI utility version

This information describes how to display help for the CLI command and the version of the XCLI utility.

About this task

The following command displays the help text for the CLI in batch mode:

```
xcli <-h | --help>
```

For details about the **help** command, see Displaying help.

The following command displays the XCLI utility version:

```
xcli <-v | --version>
```

Host and cluster management commands

This section describes the command-line interface (CLI) for host and cluster management.

- **Adding a host to a cluster**
  Use the **cluster_add_host** command to add a host to a cluster.
- **Creating a cluster**
  Use the **cluster_create** command to create a new cluster.
- **Deleting clusters**
  Use the **cluster_delete** command to delete a cluster.
- **Listing clusters**
  Use the **cluster_list** command to retrieve information about a specific cluster, or about all of them.
- **Removing a host from a cluster**
  Use the **cluster_remove_host** command to remove a host from a cluster.
- **Renaming clusters**
  Use the **cluster_rename** command to rename a cluster.
- **Adding a port to a host**
  Use the **host_add_port** command to add a port address to a host.
- **Defining a new host**
  Use the **host_define** command to define a new host to connect to the storage system.
- **Deleting a host**
  Use the **host_delete** command to delete a host.
- **Listing hosts**
  Use the **host_list** command to list a specific host or all hosts.

- **Listing ports**
  Use the **host_list_ports** command to list all the ports of a host.
- **Removing a port from a host**
  Use the **host_remove_port** command to remove a port from a host.
- **Renaming a host**
  Use the **host_rename** command to rename a host.
- **Updating a host definition**
  Use the **host_update** command to update a host definition.
- **Mapping a volume to a host or cluster**
  Use the **map_vol** command to map a volume to a host or a cluster.
- **Listing the mapping of volumes to hosts, clusters, or domains**
  Use the **mapping_list** command to list the mapping of volumes to a specified host, cluster, or domain.
- **Setting the special type of hosts or clusters**
  Use the **special_type_set** command to set the special type of a host or a cluster.
- **Listing hosts/clusters to which a volume is mapped**
  Use the **vol_mapping_list** command to list all hosts and clusters to which a volume is mapped.
- **Unmapping a volume from a host or cluster**
  Use the **unmap_vol** command to unmap a volume from a host or a cluster.
- **Set the default idle time before unmapping a volume**
  Use the **unmap_vol_set_default_idle_time** command to set the default idle time required for a volume before unmapping it.
- **Retrieving the default idle time before unmapping a volume**
  Use the **unmap_vol_get_default_idle_time** command to retrieve the default idle time required for a volume before unmapping it.
- **Creating a performance class**
  Use the **perf_class_create** command to create a performance class.
- **Deleting a performance class**
  Use the **perf_class_delete** command to delete a performance class.
- **Renaming a performance class**
  Use the **perf_class_rename** command to rename a performance class.
- **Listing details of performance classes**
  Use the **perf_class_list** command to list performance classes.
- **Adding a host to a performance class**
  Use the **perf_class_add_host** command to add a host to a performance class.
- **Removing a host from its performance class**
  Use the **perf_class_remove_host** command to remove a host from its performance class.
- **Adding a pool to a performance class**
  Use the **perf_class_add_pool** command to add a pool to a performance class.
- **Removing a pool from its performance class**
  Use the **perf_class_remove_pool** command to remove a pool from its performance class.
- **Adding a volume to a performance class**
  Use the **perf_class_add_vol** command to add a volume to a performance class.
- **Removing a volume from its performance class**
  Use the **perf_class_remove_vol** command to remove a volume from its performance class.
- **Adding a domain to a performance class**
  Use the **perf_class_add_domain** command to add a domain to a performance class.
- **Removing a domain from its performance class**
  Use the **perf_class_remove_domain** command to remove a domain from its performance class.
- **Setting the rate for a performance class**
  Use the **perf_class_set_rate** command to set the rate for a performance class.
- **Listing host profiles**
  Use the **host_profile_list** command to list all host profiles.
- **Updating the host profile**
  Use the **host_profile_set** command to update the host profile.
- **Removing the profile of the specified host**
  Use the **host_profile_clear** command to remove the profile of the specified host.
- **Enabling the host profiler**
  Use the **host_profiler_enable** command to enable the host profiler functionality.
- **Disabling the host profiler**
  Use the **host_profiler_disable** command to disable the host profiler functionality.

Adding a host to a cluster

Use the **cluster_add_host** command to add a host to a cluster.

```
cluster_add_host cluster=ClusterName host=HostName map=MapName
```

Parameters

Name	Type	Description	Mandatory
cluster	Object name	Name of the cluster to contain the host.	Y
host	Object name	Host to be added to the cluster.	Y
map	Enumeration	Defines whether to override the cluster mapping with the host mapping or vice versa, or append the cluster mapping on top of the host mapping.	Y

If the host already belongs to another cluster, the command fails. If the host already belongs to the specified cluster, the operation completes successfully, but has no effect.

Using the map parameter:

- If map=cluster, the mapping of the host and host type are overriden with the cluster's mapping and type.
- If map=host, the mapping of the cluster and its host type are overridden with the host's mapping and type. Use this value to add a host to an empty cluster, so that the cluster will obtain the host's mapping.
- If map=clusterWithHostExceptions, the host keeps its mapping and the cluster mapping is appended on top of it.

The host or cluster receives a single SCSI unit attention message, even if the change affects multiple volumes.

Example:

```
cluster_add_host cluster=Cluster1 host=Host1 map=cluster
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- HOST_BAD_NAME
  The host name does not exist.
- CLUSTER_BAD_NAME
  The cluster name does not exist.
- HOST_BELONGS_TO_ANOTHER_CLUSTER
  This host already belongs to another cluster.
- HOST_AND_CLUSTER_HAVE_CONFLICTING_MAPPINGS
  Host mapping conflicts with cluster mapping.
- HOST_AND_CLUSTER_HAVE_DIFFERENT_MAPPING_TYPE
  The host mapping type is not the same as the cluster mapping type.
- HOST_NOT_IN_CLUSTERS_DOMAINS
  The host is not part of all of the domains the cluster is attached to.

Creating a cluster

Use the **cluster_create** command to create a new cluster.

```
cluster_create cluster=ClusterName [ domain=DomainList ]
```

Parameters

Name	Type	Description	Mandatory	Default
cluster	Object name	Name of the cluster to be created.	Y	N/A
domain	N/A	The cluster will be attached to the specified domains. To define more than one domain, separate them with a comma. To attach the cluster to all existing domains, use "*".	N	none

The newly created cluster does not contain hosts, and has the default type, but no mapping.

Example:

```
cluster_create cluster=Cluster1
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- CLUSTER_NAME_EXISTS
  The cluster name already exists.
- MAX_CLUSTERS_REACHED
  The maximum number of defined clusters is already reached.
- DOMAIN_DOESNT_EXIST
  The domain does not exist.

Deleting clusters

Use the **cluster_delete** command to delete a cluster.

```
cluster_delete cluster=ClusterName
```

Parameters

Name	Type	Description	Mandatory
cluster	Object name	Cluster to be deleted.	Y

This command deletes a cluster. All hosts contained in the cluster remain active and are not deleted. The special type of each host is set to the cluster's special type. The mapping of each host is set to the cluster's mapping. No I/O interruption is caused by this command.

Example:

```
cluster_delete cluster=Cluster1
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_DELETE_ASSIGNED_CLUSTER
  Cluster *Cluster*' includes hosts. Are you sure you want to delete it?

Return codes

- CLUSTER_BAD_NAME
  The cluster name does not exist.

Listing clusters

Use the **cluster_list** command to retrieve information about a specific cluster, or about all of them.

```
cluster_list [ cluster=ClusterName ] [ domain=DomainName ]
```

Parameters

Name	Type	Description	Mandatory	Default
cluster	Object name	Name of cluster to be listed.	N	All clusters.
domain	Object name	The domain name.	N	All Domains

The output provides each cluster's special type, and comma-separated lists of hosts, users, and user groups.

Example:

```
cluster_list
```

Output:

```
Name        Hosts   Type       Creator            User Group
----------- ------- ---------- ------------------ ------------
Cluster1            default    xiv_maintenance
```

Field ID	Field output	Default position
name	Name	1
hosts	Hosts	2
type	Type	3
creator	Creator	4
user_group	User Group	5

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Removing a host from a cluster

Use the **cluster_remove_host** command to remove a host from a cluster.

```
cluster_remove_host cluster=ClusterName host=HostName
```

Parameters

Name	Type	Description	Mandatory
cluster	Object name	Cluster name.	Y
host	Object name	Host to be removed from cluster.	Y

This command removes the specified host from a cluster. The host then no longer belongs to any cluster. The host's special type and mapping remain identical to the cluster's special type and mapping, and therefore, I/O is not interrupted. The association of the host with user or user groups remains the same as the cluster's association.

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- HOST_BAD_NAME
  The host name does not exist.
- CLUSTER_BAD_NAME
  The cluster name does not exist.
- HOST_NOT_IN_CLUSTER
  This host does not belong to the specified cluster.

Renaming clusters

Use the **cluster_rename** command to rename a cluster.

```
cluster_rename cluster=ClusterName new_name=Name
```

Parameters

Name	Type	Description	Mandatory
cluster	Object name	Cluster to be renamed.	Y
new_name	Object name	New name of cluster.	Y

This command renames the specified cluster.

Example:

```
cluster_rename cluster=Cluster1 new_name=Cluster2
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- CLUSTER_BAD_NAME
  The cluster name does not exist.
- CLUSTER_NAME_EXISTS
  The cluster name already exists.

Adding a port to a host

Use the **host_add_port** command to add a port address to a host.

```
host_add_port host=HostName < fcaddress=wwpn | iscsi_name=iSCSIName [ num_of_visible_targets=num |
num_of_visible_targets_per_network=num ] >
```

Parameters

Name	Type	Description	Mandatory	Default
host	Object name	Host name.	Y	N/A
fcaddress	N/A	FC address of the added port.	N	N/A
iscsi_name	iSCSI initiator name	iSCSI initiator name of the newly added port.	N	N/A
num_of_visible_targets	Integer	Limit the maximum number of target IP addresses which will be reported on iSCSI discoveries invoked by this initiator port. Valid values: 0 (unlimited), 2-64.	N	0
num_of_visible_targets_per_network	Integer	Maximum number of target IP addresses to be reported on iSCSI discoveries invoked by this initiator port per network. Valid values: 0 (unlimited), 2-64.	N	0

The FC port address or iSCSI initiator (port) name assigned to the host must be unique per storage system. The FC port name must be exactly 16 characters long, in hexadecimal format.

Only the following alphanumeric characters are valid: 0-9, A-F, a-f. In addition to the 16 characters, colons (:) may be used as separators in the 16 character port name. The iSCSI initiator name may not exceed 253 characters and may not contain any blank spaces.

Example:

```
host_add_port host=Host1 fcaddress=5001738035C601C0
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- HOST_BAD_NAME
  The host name does not exist.
- HOST_PORT_EXISTS
  A host with this port ID is already defined.
- ISCSI_HOST_ILLEGAL_PORT_NAME
  The port name for iSCSI Host is illegal.

  Troubleshooting: Port names for iSCSI Hosts must contain only printable characters.
- MAX_PORTS_REACHED
  The maximum number of ports defined in the system is already reached.
- TARGET_PORT_BAD_ADDRESS
  The remote port address is illegal or does not belong to the remote target.
- PORT_EXISTS
  The port is already defined.
- OLVM_LINK_IS_NOT_UP
  The IBM Hyper-Scale Mobility link is not up. The mapping list cannot be updated.
- TARGET_NOT_CONNECTED
  There is currently no connection to the target system.
- REMOTE_MAX_VIRTUAL_HOSTS_REACHED
  The maximum number of defined remote virtual hosts is already reached.
- INVALID_NUM_OF_TARGETS
  The specified value representing the number of visible targets is invalid. It can be either 0 (unlimited), or a number in the range from 2 through 64.
- PARAMETER_CANNOT_BE_UPDATED_ON_SYSTEM_HAVING_IPINTERFACE_WITH_VLAN
  The num_of_visible_targets parameter cannot be used in a system with an IP interface that is configured for VLAN.

Defining a new host

Use the **host_define** command to define a new host to connect to the storage system.

```
host_define host=HostName [ cluster=ClusterName ]
```

```
[ iscsi_chap_name=iscsiChapName iscsi_chap_secret=iscsiPass ] [ domain=DomainList ]
```

Parameters

Name	Type	Description	Mandatory	Default
host	Object name	The name of the host to be created.	Y	N/A
cluster	Object name	The name of the cluster to contain the newly created host.	N	No cluster.
iscsi_chap_name	String	The host's CHAP name identifier.	N	[None]
iscsi_chap_secret	String	The password of the initiator used to authenticate to the system when CHAP is enabled.	N	[None]
domain	N/A	The domains the cluster will be attached to. To include more than one domain, separate them with a comma. To include all existing domains, use an asterisk ("*").	N	none

The name of the host must be unique in the system.

Use the host_add_port command to add port addresses to this host (see Adding a port to a host for details). Specifying the cluster is optional.

The parameters iscsi_chap_name and iscsi_chap_secret must be either both specified or both unspecified.

If iscsi_chap_secret does not conform to the required secret length (96-128 bits), the command will fail.

The command checks whether the iscsi_chap_name and iscsi_chap secret are unique. In case they are not, an error message is displayed, but the command completes.

The secret has to be between 96 bits and 128 bits. There are 3 ways to enter the secret:

- *Base64*: Requires to prefix the entry with 0b. Each subsequent character entered is treated as a 6-bit equivalent length
- *Hex*: Requires to prefix the entry with 0x. Each subsequent character entered is treated as a 4-bit equivalent length
- *String*: Requires no prefix (cannot be prefixed with 0b or 0x). Each character entered is treated as a 8 bit equivalent length

Example:

```
host_define host=server1
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- ISCSI_CHAP_NAME_AND_SECRET_NOT_UNIQUE
  Both iSCSI CHAP name and secret are already used by another host. Are you sure you want to reuse those values?

Return codes

- HOST_NAME_EXISTS
  The host name already exists.
- MAX_HOSTS_REACHED
  The maximum number of defined hosts is already reached.
- CLUSTER_BAD_NAME
  The cluster name does not exist.
- DOMAIN_DOESNT_EXIST
  The domain does not exist.
- HOST_NOT_IN_CLUSTERS_DOMAINS
  The host is not part of all of the domains the cluster is attached to.

Deleting a host

Use the **host_delete** command to delete a host.

```
host_delete host=HostName
```

Parameters

Name	Type	Description	Mandatory
host	Object name	The host name.	Y

After this command is executed, the deleted host can no longer connect to the system, and I/O requests from this host are not handled.

Example:

```
host_delete host=mailserver
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed

User Category	Permission
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_DELETE_HOST
  Are you sure you want to delete host *Host*?

Return codes

- HOST_BAD_NAME
  The host name does not exist.

Listing hosts

Use the **host_list** command to list a specific host or all hosts.

```
host_list [ host=HostName ] [ perf_class=perfClassName ] [ domain=DomainName ]
```

Parameters

Name	Type	Description	Mandatory	Default
host	Object name	The host name.	N	All hosts.
perf_class	Object name	The name of a performance class.	N	no filter.
domain	Object name	The domain name.	N	All Domains

This command lists all the hosts in the system.

A host name can be specified to list only a specific host or all the hosts.

The list contains the following comma separated information:

- Port addresses
- Containing cluster, if one exists
- Associated users and user groups

Example:

```
host_list host=mailserver
```

Output:

```
Name     Type      FC Ports   iSCSI Ports   User Group   Cluster
-------- --------- ---------- ------------- ------------ ---------
host_4   default               iscsi_4
host_5   default               iscsi_5
host_6   default               iscsi_6
host_7   default               iscsi_7
host_8   default               iscsi_8
host_9   default               iscsi_9
```

Field ID	Field output	Default position
name	Name	1
type	Type	2
fc_ports	FC Ports	3
iscsi_ports	iSCSI Ports	4
creator	Creator	N/A
user_group	User Group	5
cluster	Cluster	6
perf_class	Performance Class	7
iscsi_chap_name	iSCSI CHAP Name	N/A

Access control

User Category	Permission
Storage administrator	Allowed

User Category	Permission
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Listing ports

Use the **host_list_ports** command to list all the ports of a host.

```
host_list_ports host=HostName [ domain=DomainName ]
```

Parameters

Name	Type	Description	Mandatory	Default
host	Object name	The host name.	Y	N/A
domain	Object name	The domain name.	N	All Domains

Example:

```
host_list_ports host=tlib_host_pro125_fc0
```

Output:

```
Host                  Type        Port name
--------              ---------   ---------
tlib_host_pro125_fc0  FC          100000062B125CD0
```

Field ID	Field output	Default position
host	Host	1
type	Type	2
port_name	Port Name	3
num_of_visible_targets	iSCSI targets limit	4
num_of_visible_targets_per_network	iSCSI targets limit per network	5

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Removing a port from a host

Use the **host_remove_port** command to remove a port from a host.

```
host_remove_port host=HostName < fcaddress=wwpn | iscsi_name=iSCSIName >
```

Parameters

Name	Type	Description	Mandatory
host	Object name	The host name.	Y
fcaddress	N/A	FC address of the port to be removed.	N
iscsi_name	iSCSI initiator name	iSCSI initiator name of the port to be removed.	N

Example:

```
host_remove_port host=host1 iscsi_name=iscsi1
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- HOST_BAD_NAME
  The host name does not exist.
- PORT_DOES_NOT_BELONG_TO_HOST
  The port ID belongs to another host.
- HOST_PORT_DOES_NOT_EXIST
  The port ID is not defined.
- ISCSI_HOST_ILLEGAL_PORT_NAME
  The port name for iSCSI Host is illegal.

  Troubleshooting: Port names for iSCSI Hosts must contain only printable characters.
- OLVM_LINK_IS_NOT_UP
  The IBM Hyper-Scale Mobility link is not up. The mapping list cannot be updated.
- TARGET_NOT_CONNECTED
  There is currently no connection to the target system.
- TARGET_PORT_BAD_ADDRESS
  The remote port address is illegal or does not belong to the remote target.
- HOST_PORT_EXISTS
  A host with this port ID is already defined.
- MAX_PORTS_REACHED
  The maximum number of ports defined in the system is already reached.
- PORT_EXISTS
  The port is already defined.
- REMOTE_MAX_VIRTUAL_HOSTS_REACHED
  The maximum number of defined remote virtual hosts is already reached.

Renaming a host

Use the **host_rename** command to rename a host.

```
host_rename host=HostName new_name=Name
```

Parameters

Name	Type	Description	Mandatory
host	Object name	The original host name.	Y
new_name	Object name	The new host name. Must be unique in the system.	Y

The new host name must be unique in the system.

The command still succeeds even if the new name is identical to the current name.

Example:

```
host_rename host=server2 new_name=mailserver
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed

User Category	Permission
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- HOST_BAD_NAME
  The host name does not exist.
- HOST_NAME_EXISTS
  The host name already exists.

Updating a host definition

Use the **host_update** command to update a host definition.

```
host_update host=HostName [ iscsi_chap_name=iscsiChapName ] [ iscsi_chap_secret=iscsiPass ]
```

Parameters

Name	Type	Description	Mandatory	Default
host	Object name	Name that represents the host to the storage system.	Y	N/A
iscsi_chap_name	String	The host's CHAP name identifier	N	[unchanged]
iscsi_chap_secret	String	The password of the initiator used to authenticate to the storage system when CHAP is enabled	N	[unchanged]

The command carries out the following CHAP-related checks:

- The parameters iscsi_chap_name and iscsi_chap_secret must be either both specified or both unspecified.
  These parameters have to be unique. In case they are not, an error message is displayed, but the command completes.
- The secret needs to be between 96 bits and 128 bits. There are 3 ways to enter the secret:
  - *Base64*: Requires to prefix the entry with 0b. Each subsequent character entered is treated as a 6-bit equivalent length
  - *Hex*: Requires to prefix the entry with 0x. Each subsequent character entered is treated as a 4-bit equivalent length
  - *String*: Requires no prefix (cannot be prefixed with 0b or 0x). Each character entered is treated as an 8-bit equivalent length
- If iscsi_chap_secret does not conform with the required secret length (96-128 bits), the command fails.

Changing the iscsi_chap_name and/or iscsi_chap_secret:

- A warning message will be displayed stating that the changes will apply only next time the host is connected.

Example:

```
host_update host iscsi_chap_name iscsi_chap_secret
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- ISCSI_CHAP_NAME_AND_SECRET_NOT_UNIQUE
  Both iSCSI CHAP name and secret are already used by another host. Are you sure you want to reuse those values?
- ISCSI_CHAP_SECRET_NOT_UNIQUE
  iSCSI CHAP secret is already used by another host. Are you sure you want to reuse this value?
- ISCSI_CHAP_NAME_NOT_UNIQUE
  iSCSI CHAP name is already used by another host. Are you sure you want to reuse this value?

Return codes

- HOST_BAD_NAME
  The host name does not exist.
- ISCSI_CHAP_NAME_EMPTY
  CHAP name should be a non-empty string.
- ISCSI_CHAP_NAME_TOO_LONG
  CHAP name is too long.
- ISCSI_CHAP_SECRET_EMPTY
  CHAP secret should be a non-empty string.
- ISCSI_CHAP_SECRET_BAD_SIZE
  CHAP secret should be 12 to 16 bytes long.
- ISCSI_CHAP_SECRET_BAD_HEX_FORMAT
  CHAP secret is an illegal hexadecimal number or its size is illegal. It should be 24 to 32 hexadecimal digits.

Mapping a volume to a host or cluster

Use the **map_vol** command to map a volume to a host or a cluster.

```
map_vol <host=HostName | cluster=ClusterName> vol=VolName lun=LUN [ override=<no|yes> ]
```

Parameters

Name	Type	Description	Mandatory	Default
host	Object name	Host name.	N	N/A
cluster	Object name	Cluster name.	N	N/A
vol	Object name	Volume name.	Y	N/A
lun	Integer	LUN identifier.	Y	N/A
override	Boolean	Override the existing mapping.	N	no

This command maps a volume to a host or to a cluster. It maps the volume to all the hosts that are contained in the cluster.

The command fails if:

- The specified host is contained in a cluster, because in this case the mapping must be done through the cluster.
- Another volume is mapped to the same LUN for this cluster/host, and the override parameter is not specified.
  - If the override parameter is specified, the mapping is replaced. The host (or all hosts in the cluster) will see continuous mapping of volume to this LUN with a different content, and probably size.
- Mapping to a cluster, if the LUN was defined as an exception.
  - Whenever the LUN is defined as an exception, map it directly to the host.

Access control

User Category	Permission	Condition
Storage administrator	Allowed	N/A
Storage integration administrator	Allowed	N/A
Application administrator	Conditionally Allowed	This volume is a snapshot. The master volume of this snapshot is mapped to a host or cluster that is associated with the user executing this command. This snapshot was created by an application administrator.
Security administrator	Disallowed	N/A
Read-only users	Disallowed	N/A
Technicians	Disallowed	N/A

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_PERFORM_HOST_SPECIFIC_MAPPING
  '*Host*' is part of a cluster. Are you sure you want to map this volume only for that specific host?
- ARE_YOU_SURE_YOU_WANT_TO_MAP_VOLUME
  Are you sure you want to map volume *Volume*, which is already mapped to another host/cluster?

Return codes

- HOST_BAD_NAME
  The host name does not exist.

- HOST_BELONGS_TO_CLUSTER
  This host already belongs to a cluster.
- CLUSTER_BAD_NAME
  The cluster name does not exist.
- VOLUME_BAD_NAME
  The volume name does not exist.
- SNAPSHOT_IS_INTERNAL
  Internal snapshots cannot be mapped, modified or deleted.
- VOLUME_ALREADY_ASSIGNED
  Mapping conflict: the volume is already assigned.
- LUN_ALREADY_IN_USE
  Mapping conflict: LUN is already in use.
- EXT_LUN_ILLEGAL
  The LUN is out of range or does not exist.
- VOLUME_HAS_HOST_SPECIFIC_MAPPING
  The specified volume is currently mapped to another LUN in a host-specific mapping.
- LUN_HAS_HOST_SPECIFIC_MAPPING
  The specified LUN currently has another volume mapped in a host-specific mapping.
- VOLUME_IS_NON_PROXY_OLVM_DESTINATION
  The volume is in an IBM Hyper-Scale Mobility migration state.
- ISCSI_HOST_ILLEGAL_PORT_NAME
  The port name for iSCSI Host is illegal.

  Troubleshooting: Port names for iSCSI Hosts must contain only printable characters.
- MAX_PORTS_REACHED
  The maximum number of ports defined in the system is already reached.
- OLVM_LINK_IS_NOT_UP
  The IBM Hyper-Scale Mobility link is not up. The mapping list cannot be updated.
- HOST_PORT_EXISTS
  A host with this port ID is already defined.
- OPERATION_DENIED_OBJECT_MANAGED
  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.
- REMOTE_MAX_VIRTUAL_HOSTS_REACHED
  The maximum number of defined remote virtual hosts is already reached.
- VOLUME_HAS_INACTIVE_DATA_MIGRATION
  Cannot map a volume that has an inactive data migration.
- TARGET_NOT_CONNECTED
  There is currently no connection to the target system.
- REMOTE_TARGET_NOT_CONNECTED
  There is currently no connection from the target system.
- VOLUME_IS_AN_UNAVAILABLE_HYPERSWAP_PEER
  The operation is not permitted on a HyperSwap target which is unavailable for I/O.
- HOST_TYPE_IS_NOT_CONFIGURED
  Cannot associate a HyperSwap volume with a host of unconfigured type. IMPORTANT: Read the HyperSwap chapter in the 'Best Practices' document to understand the solution requirements.

Listing the mapping of volumes to hosts, clusters, or domains

Use the **mapping_list** command to list the mapping of volumes to a specified host, cluster, or domain.

```
mapping_list [ host=HostName | cluster=ClusterName ] [ domain=DomainList ]
```

Parameters

Name	Type	Description	Mandatory	Default
host	Object name	List of hosts to show the mapping from. To define more than one host, separate each one with a comma. To specify all existing non-internal hosts, use "*". If no host or cluster is defined, the result returns all non-internal mappings of both hosts and clusters. **Note:** Only a host or cluster argument may be provided at a time.	N	N/A

Name	Type	Description	Mandatory	Default
cluster	Object name	List of clusters to show the mapping from. To define more than one cluster, separate each one with a comma. To specify all existing non-internal clusters, use "*". If no host or cluster is defined, the result returns all non-internal mappings of both hosts and clusters. **Note:** Only a host or cluster argument may be provided at a time.	N	N/A
domain	N/A	List of domains to show the mapping from. To define more than one domain, separate them with a comma. To specify all existing domains, use "*".	N	All user domains

Field ID	Field output	Default position
lun	LUN	1
volume	Volume	2
proxy	Proxy	3
size	Size	4
master	Master	5
wwn	WWN	6
locked	Locked	7
host	Host	8

Example:

```
mapping_list host=demo__host_1,demo_host_fc10000006072d0190
```

Output:

```
LUN   Volume                                                Proxy   Size   Master
----- ------------------------------------------------- ------- ------ ------------------
0     vol-2693072-0006                                    no      103
1     vol-2693072-0007                                    no      103
2     cg-2693072-0005.snap_group_00001.vol-2693072-0006   no      103    vol-2693072-0006
3     cg-2693072-0005.snap_group_00001.vol-2693072-0007   no      103    vol-2693072-0007
4     vol-2693172-0013                                    no      103
5     vol-2693172-0013.snapshot_00001                     no      103    vol-2693172-0013
```

```
WWN                               Locked   Host
--------------------------------- -------- ------------------------------------
60017380000035c7000000000000000a  no       tlib_host_host081_fc21000024ff2c4cf7
60017380000035c7000000000000000b  no       tlib_host_host081_fc21000024ff2c4cf7
60017380000035c7000000000000000c  yes      tlib_host_host081_fc21000024ff2c4cf7
60017380000035c7000000000000000c  yes      tlib_host_host081_fc21000024ff2c4cf7
60017380000035c70000000000000011  no       tlib_host_host081_fc21000024ff2c4cf7
60017380000035c70000000000000012  no       tlib_host_host081_fc21000024ff2c4cf7
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Return codes

- HOST_BAD_NAME
  The host name does not exist.
- CLUSTER_BAD_NAME
  The cluster name does not exist.
- TOO_MANY_MAPPINGS
  There are too many mappings to display.
- DOMAIN_DOESNT_EXIST
  The domain does not exist.

Setting the special type of hosts or clusters

Use the **special_type_set** command to set the special type of a host or a cluster.

```
special_type_set <host=HostName | cluster=ClusterName>  type=<default|hpux|zvm|Linux|ESXi|Windows2008|AllOther
s>
```

Parameters

Name	Type	Description	Mandatory
host	Object name	Host name.	N
cluster	Object name	Cluster name.	N
type	Enumeration	Special map type.	Y

Make sure to define the type parameter prior to attaching HyperSwap volumes to the host. The supported special types are **hpux**, **zvm**, **Linux**, **ESXi** and **Windows2008**. For any other operating system, select **AllOthers**.

Note:
If you need to modify the type parameter, make sure to do it when creating a new host definition. Changing the type when volumes are already attached to the host, will cause loss of access to the host.

Note:
Modifying the type parameter to 'Linux' or 'ESXi' is only permitted if host or cluster are not mapped with a LUN greater than 255.

Example:

```
special_type_set host=tlib_host_pro26_fc0 type=zvm
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- HOST_BAD_NAME
  The host name does not exist.
- HOST_BELONGS_TO_CLUSTER
  This host already belongs to a cluster.
- HOST_HAS_MAPPED_VOLUMES
  The host already has mapped volumes.
- CLUSTER_BAD_NAME
  The cluster name does not exist.
- CLUSTER_HAS_MAPPED_VOLUMES
  The cluster already has mapped volumes.
- HOST_HAS_LUN_GREATER_THAN_255
  The host cannot be mapped to a LUN greater than 255 when setting type to 'Linux' or 'ESXi'.
- CLUSTER_HAS_LUN_GREATER_THAN_255
  The cluster cannot be mapped to a LUN greater than 255 when setting type to 'Linux' or 'ESXi'.

Listing hosts/clusters to which a volume is mapped

Use the **vol_mapping_list** command to list all hosts and clusters to which a volume is mapped.

```
vol_mapping_list vol=VolName
```

Parameters

Name	Type	Description	Mandatory
vol	Object name	Volume name.	Y

This command lists all the hosts and clusters to which a volume is mapped, as well as hosts that are part of a cluster and have host-specific mapping to the volume. The output list contains two columns: name of host/cluster and type (host or cluster).

Field ID	Field output	Default position
host	Host/Cluster	1
type	Type	2

Field ID	Field output	Default position
lun	LUN	3

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.

Unmapping a volume from a host or cluster

Use the **unmap_vol** command to unmap a volume from a host or a cluster.

```
unmap_vol <host=HostName | cluster=ClusterName> vol=VolName [ idle_seconds=IdleSeconds ]  [ force=<yes|no> ]
```

Parameters

Name	Type	Description	Mandatory	Default
host	Object name	Host name.	N	N/A
cluster	Object name	Cluster name.	N	N/A
vol	Object name	Volume name.	Y	N/A
idle_seconds	Integer	How many seconds the volume needs to be idle before unmapping.	N	-1
force	Boolean	Force completing the unmap operation, even if there exists in-flight I/O that has not completed on the proxy volume.	N	no

The command to unmap from a cluster will unmap the volume from all the hosts that are contained in that cluster.

The command fails if the specified host is contained in a cluster. In this case, the unmapping of the host must be performed through the cluster.

The command does not fail when the volume is not mapped to the host/cluster.

Using this command with unmap_vol_set_default_idle_time: The default value of the idle_seconds parameter can be set before running the unmap_volume command.

The command takes some time to process: If the command fails with VOLUME_NOT_IDLE (see the completion codes table below), wait one minute to allow the host to complete background writes, and try the command again.

Access control

User Category	Permission	Condition
Storage administrator	Allowed	N/A
Storage integration administrator	Allowed	N/A
Application administrator	Conditionally Allowed	The volume is a snapshot, where its master volume is mapped to a host or cluster associated with the user and the snapshot was created by an application administrator.
Security administrator	Disallowed	N/A
Read-only users	Disallowed	N/A
Technicians	Disallowed	N/A

Warnings

- HA_SLAVE_NOT_CONNECTED
  The secondary volume in this HyperSwap relation is not connected to the primary volume. Are you sure you want to unmap the volume?

  Troubleshooting: It is not recommended to unmap the secondary volume if the relation is still in use.

Return codes

- HOST_BAD_NAME
  The host name does not exist.

- HOST_BELONGS_TO_CLUSTER
  This host already belongs to a cluster.
- CLUSTER_BAD_NAME
  The cluster name does not exist.
- VOLUME_BAD_NAME
  The volume name does not exist.
- SNAPSHOT_IS_INTERNAL
  Internal snapshots cannot be mapped, modified or deleted.
- VOLUME_IS_NON_PROXY_OLVM_DESTINATION
  The volume is in an IBM Hyper-Scale Mobility migration state.
- OPERATION_DENIED_OBJECT_MANAGED
  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.
- OLVM_LINK_IS_NOT_UP
  The IBM Hyper-Scale Mobility link is not up. The mapping list cannot be updated.
- ISCSI_HOST_ILLEGAL_PORT_NAME
  The port name for iSCSI Host is illegal.

  Troubleshooting: Port names for iSCSI Hosts must contain only printable characters.
- MAX_PORTS_REACHED
  The maximum number of ports defined in the system is already reached.
- HOST_PORT_EXISTS
  A host with this port ID is already defined.
- VOLUME_NOT_IDLE
  The volume was not idle before unmapping. Check connected hosts and idle timeout.
- MAPPING_IS_NOT_DEFINED
  The requested mapping is not defined.
- REMOTE_MAX_VIRTUAL_HOSTS_REACHED
  The maximum number of defined remote virtual hosts is already reached.
- TARGET_NOT_CONNECTED
  There is currently no connection to the target system.
- REMOTE_TARGET_NOT_CONNECTED
  There is currently no connection from the target system.

Set the default idle time before unmapping a volume

Use the **unmap_vol_set_default_idle_time** command to set the default idle time required for a volume before unmapping it.

```
unmap_vol_set_default_idle_time idle_time_seconds=IdleSeconds
```

Parameters

Name	Type	Description	Mandatory
idle_time_seconds	Integer	Defines how many seconds the volume needs to be idle before unmapping.	Y

Example:

```
unmap_vol_set_default_idle_time idle_time_seconds=10
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Retrieving the default idle time before unmapping a volume

Use the **unmap_vol_get_default_idle_time** command to retrieve the default idle time required for a volume before unmapping it.

```
unmap_vol_get_default_idle_time
```

Example:

```
unmap_vol_get_default_idle_time
```

Output:

```
idle_time_seconds = "0"
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Creating a performance class

Use the **perf_class_create** command to create a performance class.

```
perf_class_create perf_class=perfClassName [ type=<shared|independent> ]
```

Parameters

Name	Type	Description	Mandatory	Default
perf_class	String	The name of a performance class.	Y	N/A
type	Enumeration	Determines if associated objects will be limited independently or share the same limit.	N	shared

The performance class name must be unique. Up to 1000 classes can be created.

Example:

```
perf_class_create perf_class=p1
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- PERF_CLASS_EXISTS
  The performance class already exists.
- MAX_PERF_CLASSES_REACHED
  The maximum number of defined performance classes is already reached.

Deleting a performance class

Use the **perf_class_delete** command to delete a performance class.

```
perf_class_delete perf_class=perfClassName
```

Parameters

Name	Type	Description	Mandatory
perf_class	Object name	Name of a performance class.	Y

Example:

```
perf_class_delete perf_class=p1
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_DELETE_A_PERF_CLASS
  Are you sure you want to delete performance class *Performance Class*?
- ARE_YOU_SURE_YOU_WANT_TO_DELETE_A_PERF_CLASS_ASSOCIATED_WITH_HOSTS
  Deleting performance class *Performance Class* will remove the performance limits set for hosts associated with the performance class. Are you sure you want to delete performance class *Performance Class*?
- ARE_YOU_SURE_YOU_WANT_TO_DELETE_A_PERF_CLASS_ASSOCIATED_WITH_POOLS
  Deleting performance class *Performance Class* will remove the performance limits set for pools associated with the performance class. Are you sure you want to delete performance class *Performance Class*?
- ARE_YOU_SURE_YOU_WANT_TO_DELETE_A_PERF_CLASS_ASSOCIATED_WITH_VOLUMES
  Deleting performance class *Performance Class* will remove the performance limits set for volumes associated with the performance class. Are you sure you want to delete performance class *Performance Class*?
- ARE_YOU_SURE_YOU_WANT_TO_DELETE_A_PERF_CLASS_ASSOCIATED_WITH_DOMAINS
  Deleting performance class *Performance Class* will remove the performance limits set for domains associated with the performance class. Are you sure you want to delete performance class *Performance Class*?

Return codes

- PERF_CLASS_BAD_NAME
  The performance class does not exist.

Renaming a performance class

Use the **perf_class_rename** command to rename a performance class.

```
perf_class_rename perf_class=perfClassName new_name=Name
```

Parameters

Name	Type	Description	Mandatory
perf_class	Object name	The name of an existing performance class.	Y
new_name	String	The new name for the performance class. The class new name must be unique.	Y

Example:

```
perf_class_rename perf_class=p1 new_name=perf1
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed

User Category	Permission
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- PERF_CLASS_EXISTS
  The performance class already exists.
- PERF_CLASS_BAD_NAME
  The performance class does not exist.

Listing details of performance classes

Use the **perf_class_list** command to list performance classes.

```
perf_class_list [ perf_class=perfClassName ]
```

Parameters

Name	Type	Description	Mandatory	Default
perf_class	String	Name of a performance class. If left unspecified, all performance classes will be listed.	N	All performance classes.

Field ID	Field output	Default position
name	Performance class	1
type	Class type	2
max_iops	Max IO rate(IOPS)	3
max_bw	Max BW rate(MB/sec)	4

Example:

```
perf_class_list
```

Output:

```
Performance class   Max IO rate(IOPS)   Max BW rate(MB/sec)
perf1               0                    0
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Adding a host to a performance class

Use the **perf_class_add_host** command to add a host to a performance class.

```
perf_class_add_host perf_class=perfClassName host=HostName
```

Parameters

Name	Type	Description	Mandatory
perf_class	Object name	The name of a performance class.	Y
host	Object name	The name of the host to be added to the performance class.	Y

If the host is already associated with another performance class, it will be removed from that performance class.

Example:

```
perf_class_add_host perf_class=p1 host=h1
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- PERF_CLASS_BAD_NAME
  The performance class does not exist.
- PERF_CLASS_ASSOCIATED_WITH_POOLS_OR_DOMAINS
  Performance class *Performance Class* is already being used by a pool or domain.
- PERF_CLASS_ASSOCIATED_WITH_VOLUMES
  Performance class *Performance Class* is already being used by a volume.
- HOST_BAD_NAME
  The host name does not exist.
- HOST_ALREADY_IN_PERF_CLASS
  Host *host* is already in performance class *Performance Class*.

Removing a host from its performance class

Use the **perf_class_remove_host** command to remove a host from its performance class.

```
perf_class_remove_host host=HostName
```

Parameters

Name	Type	Description	Mandatory
host	Object name	The name of the host to be removed from its performance class.	Y

Example:

```
perf_class_remove_host host=h1
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- HOST_BAD_NAME
  The host name does not exist.
- PERF_CLASS_DOES_NOT_CONTAIN_ANY_HOSTS
  The performance class is already empty.

Adding a pool to a performance class

Use the **perf_class_add_pool** command to add a pool to a performance class.

```
perf_class_add_pool perf_class=perfClassName pool=PoolName
```

Parameters

Name	Type	Description	Mandatory
perf_class	Object name	Name of a performance class	Y
pool	Object name	Name of a pool that will be added to the performance class	Y

If the pool is already associated with another performance class, it will be removed from it.

Example:

```
perf_class_add_pool perf_class=p1 pool=h1
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- PERF_CLASS_BAD_NAME
  The performance class does not exist.
- PERF_CLASS_ASSOCIATED_WITH_HOSTS
  Performance class *Performance Class* is already being used by a host.
- PERF_CLASS_ASSOCIATED_WITH_VOLUMES
  Performance class *Performance Class* is already being used by a volume.
- POOL_DOES_NOT_EXIST
  The storage pool does not exist.
- POOL_ALREADY_IN_PERF_CLASS
  Pool *pool name* is already in performance class *Performance Class*.

Removing a pool from its performance class

Use the **perf_class_remove_pool** command to remove a pool from its performance class.

```
perf_class_remove_pool pool=PoolName
```

Parameters

Name	Type	Description	Mandatory
pool	Object name	The name of the pool to be removed from its performance class.	Y

Example:

```
perf_class_remove_pool pool=h1
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed

User Category	Permission
Technicians	Disallowed

Return codes

- POOL_DOES_NOT_EXIST
  The storage pool does not exist.
- POOL_NOT_CONNECTED_TO_ANY_PERF_CLASS
  The pool is not connected to any performance class.

Adding a volume to a performance class

Use the **perf_class_add_vol** command to add a volume to a performance class.

```
perf_class_add_vol perf_class=perfClassName vol=VolName
```

Parameters

Name	Type	Description	Mandatory
perf_class	Object name	The name of a performance class.	Y
vol	Object name	The name of the volume to be added to the performance class.	Y

If the volume is already associated with another performance class, it will be removed from that.

Example:

```
perf_class_add_vol perf_class=p1 volume=h1
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- PERF_CLASS_BAD_NAME
  The performance class does not exist.
- PERF_CLASS_ASSOCIATED_WITH_HOSTS
  Performance class *Performance Class* is already being used by a host.
- PERF_CLASS_ASSOCIATED_WITH_POOLS_OR_DOMAINS
  Performance class *Performance Class* is already being used by a pool or domain.
- VOLUME_BAD_NAME
  The volume name does not exist.
- VOLUME_ALREADY_IN_PERF_CLASS
  Volume *volume name* is already in performance class *Performance Class*.

Removing a volume from its performance class

Use the **perf_class_remove_vol** command to remove a volume from its performance class.

```
perf_class_remove_vol vol=VolName
```

Parameters

Name	Type	Description	Mandatory
vol	Object name	The name of a volume to be removed from its performance class.	Y

Example:

```
perf_class_remove_vol volume=h1
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.
- VOLUME_NOT_CONNECTED_TO_ANY_PERF_CLASS
  The volume is not connected to any performance class.

Adding a domain to a performance class

Use the **perf_class_add_domain** command to add a domain to a performance class.

```
perf_class_add_domain perf_class=perfClassName domain=DomainName
```

Parameters

Name	Type	Description	Mandatory
domain	Object name	The name of the domain to be added to the performance class.	Y
perf_class	Object name	The name of a performance class.	Y

Example:

```
perf_class_add_domain perf_class=perf1 domain=d1
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- DOMAIN_DOESNT_EXIST
  The domain does not exist.
- PERF_CLASS_BAD_NAME
  The performance class does not exist.
- PERF_CLASS_ASSOCIATED_WITH_HOSTS
  Performance class *Performance Class* is already being used by a host.
- PERF_CLASS_ASSOCIATED_WITH_VOLUMES
  Performance class *Performance Class* is already being used by a volume.
- DOMAIN_ALREADY_IN_PERF_CLASS
  Domain *domain name* is already in performance class *Performance Class*.

Removing a domain from its performance class

Use the **perf_class_remove_domain** command to remove a domain from its performance class.

```
perf_class_remove_domain domain=DomainName
```

Parameters

Name	Type	Description	Mandatory
domain	Object name	The name of the domain to be removed from its performance class.	Y

Example:

```
perf_class_remove_domain domain=d1
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- DOMAIN_DOESNT_EXIST
  The domain does not exist.
- PERF_CLASS_NOT_ASSOC_WITH_DOMAIN
  The domain is not in any performance class.

Setting the rate for a performance class

Use the **perf_class_set_rate** command to set the rate for a performance class.

```
perf_class_set_rate perf_class=perfClassName [ max_io_rate=iops ] [ max_bw_rate=bw ]
```

Parameters

Name	Type	Description	Mandatory	Default
perf_class	Object name	Name of a performance class.	Y	N/A
max_io_rate	Positive integer	Specifies the performance class maximum rate in IOPS per interface module (IOPS). The maximum setting allowed is 100,000. If zero is specified, the IOPS rate will not be limited.	N	Keep unchanged.
max_bw_rate	Positive integer	Specifies the performance class maximum rate per interface module (MB/sec). The maximum setting allowed is 10,000. If zero is specified, the bandwidth rate will not be limited.	N	Keep unchanged.

Either max_io_rate, or max_bw_rate, or both must be set.

The specified rate is applied to each interface module. To calculate the limit per system, multiply the specified rate by the number of interface modules.

Example:

```
perf_class_set_rate perf_class=p1 max_io_rate=1000
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed

User Category	Permission
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- PERF_CLASS_BAD_NAME
  The performance class does not exist.
- PERF_CLASS_INVALID_RATE
  The rate set for the performance class is invalid.

Listing host profiles

Use the **host_profile_list** command to list all host profiles.

```
host_profile_list [ host=HostName ] [ domain=DomainName ]
```

Parameters

Name	Type	Description	Mandatory	Default
host	Object name	The name of the specific host whose profiles should be listed	N	>All Host Profiles.
domain	Object name	The domain name.	N	All Domains

The command lists all host profiles or a specific one.

Field ID	Field output	Default position
host_name	Host Name	1
update_time	Update Time	2
profile	Profile	3

Example:

```
host_profile_list host
```

Output:

```
Host Name                    Update Time          Profile
---------------------------- -------------------- ----------------------
host1                        2012-05-09 22:54:36  Windows 7
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Updating the host profile

Use the **host_profile_set** command to update the host profile.

```
host_profile_set profile_value=Profile
```

Parameters

Name	Type	Description	Mandatory
profile_value	String	The host profile value length up to 1024 characters	Y

Example:

```
host_profile_set profile_value=Profile
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Disallowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- HOST_PROFILE_UPDATE_TOO_FREQUENT
  Host Profile has been set too often.

  Troubleshooting: Try again after the minimal update interval time.
- HOST_BAD_NAME
  The host name does not exist.
- MAX_HOST_PROFILES_REACHED
  The maximum number of defined host profiles is already reached.

Removing the profile of the specified host

Use the **host_profile_clear** command to remove the profile of the specified host.

```
host_profile_clear  host=HostName
```

Parameters

Name	Type	Description	Mandatory
host	Object name	The host name.	Y

Example:

```
host_profile_clear host
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Allowed

Return codes

- HOST_BAD_NAME
  The host name does not exist.
- HOST_PROFILE_DOES_NOT_EXIST
  No profile is defined for the requested host.

Enabling the host profiler

Use the **host_profiler_enable** command to enable the host profiler functionality.

```
host_profiler_enable
```

Example:

```
host_profiler_enable
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Disabling the host profiler

Use the **host_profiler_disable** command to disable the host profiler functionality.

```
host_profiler_disable
```

Example:

```
host_profiler_disable
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Volume management commands

This section describes the command-line interface (CLI) for volume management.

See also:

- Volume snapshot management commands
- Consistency group management commands
- Storage pool management commands

- **Clearing reservations of a volume**
  Use the **reservation_clear** command to clear reservations of a volume.
- **Listing reservation keys**
  Use the **reservation_key_list** command to list reservation keys.
- **Listing volume reservations**
  Use the **reservation_list** command to list volume reservations.
- **Copying volumes**
  Use the **vol_copy** command to copy a source volume onto a target volume.
- **Creating a volume**
  Use the **vol_create** command to create a new volume.
- **Deleting a volume**
  Use the **vol_delete** command to delete a volume.
- **Formatting a volume**
  Use the **vol_format** command to formats a volume.
- **Listing volumes**
  Use the **vol_list** command to list all volumes or a specific one.
- **Listing a volume's extended attributes**
  Use the **vol_list_extended** command to return the attributes of the volume which are not returned by **vol_list**.
- **Locking a volume**
  Use the **vol_lock** command to lock a volume so that it is read-only.
- **Renaming a volume**
  Use the **vol_rename** command to rename a volume.
- **Resizing a volume**
  Use the **vol_resize** command to resize a volume.
- **Unlocking a volume**
  Use the **vol_unlock** command to unlock a volume, so that it is no longer read-only and can be written to.

Clearing reservations of a volume

Use the **reservation_clear** command to clear reservations of a volume.

```
reservation_clear  vol=VolName
```

Parameters

Name	Type	Description	Mandatory
vol	Object name	The name of the volume to clear reservations of.	Y

Example:

```
reservation_clear vol=Vol1
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.
- LOCAL_PEER_IS_NOT_MASTER
  The local peer is not primary.

Listing reservation keys

Use the **reservation_key_list** command to list reservation keys.

```
reservation_key_list  [ vol=VolName ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	The name of the volume whose reservation keys are to be listed.	N	All volumes.

Example:

```
reservation_key_list vol=Vol2
```

Output:

```
Initiator Port     Volume Name         Reservation Key
------------------ ------------------- ------------------
100000062B151C3C   vol-dmathies-0a7    2
100000062B151C3C   vol-dobratz-23a     3
```

Field ID	Field output	Default position
initiator_port	Initiator Port	1
initiator_port_isid	Initiator ISID	2
vol_name	Volume Name	3
reg_key	Reservation Key	4

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed

User Category	Permission
Technicians	Disallowed

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.

Listing volume reservations

Use the **reservation_list** command to list volume reservations.

```
reservation_list  [ vol=VolName ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	The name of the volume whose reservations are to be listed.	N	All volumes.

Example:

```
reservation_list vol=Vol1
```

Output:

```
Volume Name   Reserving Port   Reservation Type   Persistent
vol1                           none               none

cont:
Reservation Type   Persistent Access Type   Initiator UID   PR Generation
                   none                    -1              0
```

Field ID	Field output	Description	Default position
name	Volume Name	N/A	1
reserved_by_port	Reserving Port	N/A	2
reserved_by_port_isid	Reserving ISID	N/A	3
reservation_type	Reservation Type	N/A	4
persistent_reservation_type	Persistent Reservation Type	N/A	5
access_type	Persistent Access Type	N/A	6
reserving_initiator_uid	Initiator UID	uid of reserving host	7
pr_generation	PR Generation	N/A	8
reservation_age	Reservation Age	N/A	9

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.

Copying volumes

Use the **vol_copy** command to copy a source volume onto a target volume.

```
vol_copy vol_src=VolName vol_trg=VolName
```

Parameters

Name	Type	Description	Mandatory
vol_src	Object name	Name of the source volume from which the data is to be taken.	Y
vol_trg	Object name	Name of the target volume to which the data is to be copied.	Y

This command copies a source volume onto a target volume.

All data stored on the target volume is lost and cannot be restored.

This command performs the following as a single atomic action:

- Deletes the target volume.
- Creates a new volume with the same name as the target volume and the same size as the source volume.
- Instantly copies the source volume data onto the target volume.

All volume-to-host mappings of the target volume remain intact during this process. Except for its size, the target volume retains all of its properties, including its name, ID, lock state, creation time and all other attributes.

Immediately after the completion of the command, the volumes are independent of each other and are valid for any further operations (including deletion).

If the target volume is larger then the source volume, excess storage space is freed and returned to the target volume's storage pool. If the target volume is smaller than the source volume, all storage space that is needed to support the additional volume's capacity is reserved from the storage pool.

The command fails in the following cases:

- The target is not formatted.
- The source volume is larger than the target volume, and there is not enough free space in the Storage Pool that contains the target for target volume resizing.
- The target volume has a snapshot associated with it or if the target volume is a snapshot.
- The target volume is locked.
- The target volume is part of any mirroring definitions (either master or slave).
- The source volume is a slave of a synchronous mirroring, and it is currently inconsistent due to either a re-synchronization or an initialization process.
- There is not enough free space in the Storage Pool that contains the target.

In the following example, the -y option suppresses the ARE_YOU_SURE_YOU_WANT_TO_COPY_VOLUME Y/N prompt.

Example:

```
vol_copy vol_src=DBVolume vol_trg=DBVolumeCopy
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_COPY_VOLUME
  Are you sure you want to copy the contents of volume *source Volume* to volume *target Volume*?

Return codes

- NOT_ENOUGH_SPACE
  No space to allocate for the volume's current usage.
- SOURCE_VOLUME_BAD_NAME
  The source volume name does not exist.
- SOURCE_VOLUME_DATA_MIGRATION_UNSYNCHRONIZED
  Data Migration to source volume has not completed.
- TARGET_VOLUME_BAD_NAME
  The target volume name does not exist.
- TARGET_VOLUME_LOCKED
  The target volume is locked.
- TARGET_VOLUME_HAS_MIRROR
  A mirror is defined for the target volume.
- TARGET_VOLUME_HAS_DATA_MIGRATION
  Data Migration is defined for the target volume.

- VOLUME_IS_SNAPSHOT
  The operation is not permitted on snapshots.
- VOLUME_IDENTICAL
  The same volume is defined as source and target.
- VOLUME_HAS_SNAPSHOTS
  The volume has snapshots.
- VOLUME_IS_NOT_CONSISTENT_SLAVE
  The operation not allowed on an inconsistent secondary volume.
- VOLUME_IS_NOT_CONSISTENT_OLVM_DESTINATION
  The operation not allowed on an inconsistent IBM Hyper-Scale Mobility volume.
- TARGET_VOLUME_NOT_FORMATTED
  The target volume is not formatted.
- SNAPSHOT_IS_FORMATTED
  The snapshot is formatted.
- VOLUME_TOO_BIG_TO_COPY
  The volume is too large to be copied.
- TARGET_VOLUME_HAS_OLVM
  This target volume is part of an IBM Hyper-Scale Mobility relationship.
- VOLUME_IS_OLVM_PROXY
  The volume is in an IBM Hyper-Scale Mobility Proxy phase.
- OPERATION_DENIED_OBJECT_MANAGED
  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support
- SYSTEM_OUT_OF_PHYSICAL_SPACE
  The operation not allowed while the system is out of physical space.
- VOLUME_TOO_BIG
  No space to allocate to the volume.
- VOLUME_HAS_HA
  This operation is forbidden on a volume with a HyperSwap relation.
- TARGET_VOLUME_HAS_HA
  This operation is forbidden, if the target volume is a peer in a HyperSwap relation.

Creating a volume

Use the **vol_create** command to create a new volume.

```
vol_create vol=VolName < size=GB | size_blocks=BLOCKS > pool=PoolName  [ ext_id=Identifier ]
 [ perf_class=perfClassName ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	Volume name.	Y	N/A
size	Positive integer	Volume size in GB.	N	N/A
size_blocks	Positive integer	Size in number of blocks.	N	N/A
pool	Object name	The name of the storage pool to which the volume belongs.	Y	N/A
ext_id	String	External identifier of the volume.	N	N/A
perf_class	Object name	Name of the performance class for the volume.	N	No performance class

This command is used to create a new volume. The name of the volume must be unique in the system.

The space for the volume is allocated from the specified storage pool and the volume belongs to that storage pool. Specifying the storage pool is mandatory.

When creating a volume, the storage space that is needed to support the volume's capacity is reserved from the capacity of the storage pool for the volume. The command fails if the reservation cannot be committed.

Volumes are created in increments of approximately 1 GB. In some cases, rounding of up to 5% of the total volume size can take place in order to improve internal accounting. The volume size is the actual "net" storage space, as seen by the user's applications, not including any internal overhead, such as rounding.

The volume is logically formatted at the creation time, which means that any read operation results in returning all zeros as a response.

Upon successful completion of the command, its lock state is unlocked, meaning that write, format and resize operations are allowed.

The creation time of the volume is set to the current time and is never changed.

Example:

```
vol_create vol=DBVolume size=2000 pool=DBPool
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- VOLUME_SIZE_VERY_LARGE_ARE_YOU_SURE
  The volume size is very large. It may not be possible to mirror this volume to older versions of the storage system. Are you sure?

Return codes

- VOLUME_CANNOT_HAVE_ZERO_SIZE
  The volume size cannot be zero.
- POOL_DOES_NOT_EXIST
  The storage pool does not exist.
- VOLUME_EXISTS
  The volume name already exists.
- VOLUME_BAD_PREFIX
  The volume name has a reserved prefix.
- VOLUME_TOO_BIG
  No space to allocate to the volume.
- MAX_VOLUMES_REACHED
  The maximum allowed number of volumes is already reached.
- ELECTRONIC_LICENSE_NOT_APPROVED
  Operation blocked until Electronic license approval

  Troubleshooting: Please retrieve Electronic license version and accept it
- VOLUME_SIZE_ABOVE_LIMIT
  The specified volume size is above the limit.
- INVALID_SLICE_OFFSET
  Slice offset is illegal.
- OPERATION_DENIED_OBJECT_MANAGED
  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.
- ENCRYPTION_IN_PROGRESS
  The system is in the process of changing the encryption activation state.
- DOMAIN_MAX_VOLUMES_REACHED
  The domain exceeds the maximum allowed number of volumes.
- PERF_CLASS_ASSOCIATED_WITH_HOSTS
  Performance class *Performance Class* is already being used by a host.
- PERF_CLASS_ASSOCIATED_WITH_POOLS_OR_DOMAINS
  Performance class *Performance Class* is already being used by a pool or domain.
- PERF_CLASS_BAD_NAME
  The performance class does not exist.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

Troubleshooting: Contact IBM Support

- SYSTEM_OUT_OF_PHYSICAL_SPACE
  The operation not allowed while the system is out of physical space.

Deleting a volume

Use the **vol_delete** command to delete a volume.

```
vol_delete vol=VolName
```

Parameters

Name	Type	Description	Mandatory
vol	Object name	Name of the volume to delete.	Y

After deletion, all data stored on the volume is lost and cannot be restored.

This command cannot be applied to a snapshot. To delete a snapshot, use Deleting a snapshot.

All storage space allocated (or reserved) for the volume is freed and returned to the volume's storage pool.

The volume is removed from all LUN Maps that contain a mapping of the volume.

This command deletes all snapshots associated with this volume. Even snapshots that are part of a snapshot group (this can happen when the volume was in a consistency group and was removed from it prior to the deletion).

This command cannot be applied to a volume that is part of a consistency group or to a volume that is mapped to a host or cluster.

The command succeeds regardless of the volume's lock state.

Example:

```
vol_delete vol=DBVolumeCopy
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_DELETE_VOLUME
  Are you sure you want to delete volume *Volume*?
- ARE_YOU_SURE_YOU_WANT_TO_DELETE_VOLUME_WITH_SNAPSHOTS
  Volume *Volume* has snapshots! Are you sure you want to delete this volume AND all of its snapshots?

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.
- VOLUME_HAS_MIRROR
  A mirror is defined for this volume.
- SNAPSHOT_IS_PART_OF_SNAPSHOT_GROUP
  The snapshot is part of a snapshot group.
- SNAPSHOT_IS_INTERNAL
  Internal snapshots cannot be mapped, modified or deleted.
- VOLUME_BELONGS_TO_CG

The volume belongs to a consistency group.

- VOLUME_IS_MAPPED
  The volume mapped to a host cannot be deleted.
- VOLUME_IS_BOUND
  The volume is bound to an ALU.

  Troubleshooting: Unbind the volume from the ALU.
- VOLUME_HAS_MAPPED_SNAPSHOT
  A volume with a snapshot that is mapped to a host cannot be deleted.
- SNAPSHOT_HAS_ACTIVE_SYNC_JOB
  The snapshot is currently the target of an active sync job.

  Troubleshooting: Please wait for the sync job to complete.
- SNAPSHOT_IS_CONSISTENT_ELCS
  If a mirrored volume is not consistent, then its ELCS is protected and cannot be deleted.
- VOLUME_HAS_OLVM
  An IBM Hyper-Scale Mobility relationship is defined for this volume.
- VOLUME_IS_OLVM_PROXY
  The volume is in an IBM Hyper-Scale Mobility Proxy phase.
- VOLUME_IS_OLVM_DESTINATION
  The volume is defined as an IBM Hyper-Scale Mobility destination.
- OPERATION_DENIED_OBJECT_MANAGED
  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.
- VOLUME_HAS_DATA_MIGRATION
  Data Migration is defined for this volume.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support
- VOLUME_HAS_HA
  This operation is forbidden on a volume with a HyperSwap relation.

Formatting a volume

Use the **vol_format** command to formats a volume.

```
vol_format vol=VolName
```

Parameters

Name	Type	Description	Mandatory
vol	Object name	Name of the volume to be formatted.	Y

A formatted volume returns zeros as a response to any read command.

All data stored on the volume is lost and cannot be restored.

The formatting of the volume is done logically and no data is actually written to the physical storage space allocated for the volume. This allows the command to complete instantly.

The volume's lock state must be unlocked when the command is issued.

This command fails if the volume has snapshots associated with it, or if the volume is a snapshot, or if the volume is part of any mirroring or data migration definition.

Example:

```
vol_format vol=DBVolume
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed

User Category	Permission
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_FORMAT_VOLUME
  Volume *Volume* may contain data. Formatting it will cause data loss. Are you sure you want to format volume *Volume*?

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.
- VOLUME_IS_OLVM_PROXY
  The volume is in an IBM Hyper-Scale Mobility Proxy phase.
- VOLUME_HAS_SNAPSHOTS
  The volume has snapshots.
- VOLUME_IS_SNAPSHOT
  The operation is not permitted on snapshots.
- VOLUME_LOCKED
  The volume is locked.
- VOLUME_HAS_MIRROR
  A mirror is defined for this volume.
- VOLUME_HAS_OLVM
  An IBM Hyper-Scale Mobility relationship is defined for this volume.
- VOLUME_HAS_DATA_MIGRATION
  Data Migration is defined for this volume.
- OPERATION_DENIED_OBJECT_MANAGED
  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support
- VOLUME_HAS_HA
  This operation is forbidden on a volume with a HyperSwap relation.

Listing volumes

Use the **vol_list** command to list all volumes or a specific one.

```
vol_list [ vol=VolName | pool=PoolName | cg=cgName ] [ show_proxy=<yes|no> ] [ managed=<yes|no|all> ]
[ domain=DomainName ] [ wwn=WWNString ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	Name of a specific volume to be listed.	N	All volumes.
pool	Object name	Name of a specific pool whose volumes are to be listed.	N	Volumes in all Pools.
cg	Object name	List all the volumes in this consistency group.	N	All Consistency Groups.
show_proxy	Boolean	Returns data on proxy volumes (volumes in Proxy state) as well.	N	No
managed	Boolean	Filter only volumes that are or are not managed.	N	no.
domain	Object name	The domain name.	N	All Domains
wwn	String	The WWN in string format.	N	""

This command lists volumes according to:

- Volume name
- Pool
- Consistency Group
- WWN

If no parameter is indicated, the command lists all the available volumes. In addition, the command indicates whether the volume is mirrored. This command displays the following VAAI fields (available in the XML output format):

- enable_VAAI
- user_disabled_VAAI

This command displays the following snapshot format field (available in the XML output format):

- snapshot_format

Field ID	Field output	Default position
name	Name	1
size	Size (GB)	2
size_MiB	Size (MiB)	N/A
vol_copy_type	Copy type	3
master_name	Master Name	4
cg_name	Consistency Group	5
pool_name	Pool	6
creator	Creator	7
written	Written (GB)	8
written_MiB	Written (MiB)	N/A
proxy	Proxy	N/A
capacity	Capacity (blocks)	N/A
modified	Modified	N/A
sg_name	Snapshot Group Name	N/A
delete_priority	Deletion Priority	N/A
locked	Locked	N/A
snapshot_time	Snapshot Creation Time	N/A
snapshot_time_on_master	Master Copy Creation Time	N/A
snapshot_internal_role	Snapshot Internal Role	N/A
snapshot_of	Snapshot of	N/A
sg_snapshot_of	Snapshot of Snap Group	N/A
wwn	WWN	N/A
mirrored	Mirrored	N/A
locked_by_pool	Locked by Pool	N/A
capacity_used_by_snapshots_MiB	Capacity Used by Snapshots (MiB)	N/A
short_lived_io	Short Live IO	N/A
enable_VAAI	VAAI enabled	N/A
user_disabled_VAAI	VAAI disabled by user	N/A
snapshot_format	Snapshot Format	N/A
unmap_support	Unmap Support	N/A
managed	Managed	N/A
marked	Marked	N/A
perf_class	Performance Class Name	N/A
thin_provisioning_savings	Thin Provisioning Savings (%)	N/A
est_compression_factor	Est. Compression Factor	N/A
unique_stored_data	Unique Stored Data (GB)	N/A
ha	HA Relation	N/A
target_port_group_id	TPG ID	N/A
target_port_group_state	TPG State	N/A
lock_modes	Lock Modes	N/A
copy_master_wwn	Copy Master	N/A

Example:

```
vol_list
```

Output:

```
Name                               Size (GB)   Copy Type   Master Name
---------------------------------- ----------- ----------- ------------------
DBLog                              3006        None
Dev                                2010        Copy
Dev.snapshot_00001                 2010        Snapshot    Dev
VM_1                               2010        Replicated

cont:
Consistency Group   Pool              Creator           Written (GB)
------------------- ----------------- ----------------- --------------
                    MainPool          admin             21
                    MainPool          admin             13
                    MainPool          admin             0
                    MainPool          admin             20
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Allowed
Read-only users	Allowed
Technicians	Disallowed

Listing a volume's extended attributes

Use the **vol_list_extended** command to return the attributes of the volume which are not returned by **vol_list**.

```
vol_list_extended [ vol=VolName ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	Name of a specific volume to be listed.	N	All volumes.

This command lists extended attributes of volumes according to:

- Volume name

Field ID	Field output	Default position
name	Name	1
wwn	WWN	2
product_serial_number	Product Serial Number	3
uid	UID	N/A

Example:

```
vol_list_extended
```

Output:

```
Name                WWN                                Product Serial Number
------------------- ---------------------------------- -------------------------
DBLog               60017380000035C3000000000000000A   MN035C300000000000000A
Dev                 60017380000035C3000000000000000B   MN035C300000000000000B
Dev.snapshot_00001  60017380000035C3000000000000000D   MN035C300000000000000D
Dev.snapshot_00002  60017380000035C3000000000000000E   MN035C300000000000000E
Dev.snapshot_00003  60017380000035C3000000000000000F   MN035C300000000000000F
Marketing           60017380000035C3000000000000000C   MN035C300000000000000C
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Allowed
Read-only users	Allowed
Technicians	Disallowed

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.

Locking a volume

Use the **vol_lock** command to lock a volume so that it is read-only.

```
vol_lock vol=VolName
```

Parameters

Name	Type	Description	Mandatory
vol	Object name	Name of the volume to lock.	Y

This command locks a volume so that hosts cannot write to it.

A volume that is locked is write-protected, so that hosts can read the data stored on it, but cannot change it. In addition, a locked volume cannot be formatted or resized. In general, locking a volume prevents any operation (other than deletion) that changes the volume's image.

This command succeeds when the volume's lock state is already set to the one the user is trying to apply. In this case, the lock state remains unchanged.

The lock state of a master volume is set to *unlocked* when a master volume is created.

The lock state of a snapshot is set to *locked* when a snapshot is created.

In addition to the lock state, snapshots also have a modification state. The modification state is a read-only state (which cannot be changed by the user explicitly) and it is initially set to *unmodified* when the snapshot is created. The first time a snapshot lock state is set to *unlocked*, the modification state of the snapshot is changed to *modified*, and it is never changed thereafter.

If applied on a volume that is part of an IBM Hyper-Scale Mobility relation, the command has to be acknowledged by both source and destination volumes. Otherwise, a return code is returned (see below).

Example:

```
vol_lock vol=DBVolume
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission	Condition
Storage administrator	Allowed	N/A
Storage integration administrator	Allowed	N/A
Application administrator	Conditionally Allowed	The volume is a snapshot, where its master volume is mapped to a host or cluster associated with the user and the snapshot was created by an application administrator.
Security administrator	Disallowed	N/A
Read-only users	Disallowed	N/A
Technicians	Disallowed	N/A

Return codes

- SNAPSHOT_IS_INTERNAL
  Internal snapshots cannot be mapped, modified or deleted.
- VOLUME_BAD_NAME
  The volume name does not exist.
- VOLUME_IS_OLVM_PROXY
  The volume is in an IBM Hyper-Scale Mobility Proxy phase.
- VOLUME_IS_SLAVE
  The volume is defined as a secondary volume.
- VOLUME_IS_OLVM_DESTINATION
  The volume is defined as an IBM Hyper-Scale Mobility destination.
- SNAPSHOT_IS_PART_OF_SNAPSHOT_GROUP
  The snapshot is part of a snapshot group.
- OPERATION_DENIED_OBJECT_MANAGED
  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support
- LOCAL_PEER_IS_NOT_MASTER
  The local peer is not primary.
- REMOTE_TARGET_NOT_CONNECTED
  There is currently no connection from the target system.
- TARGET_NOT_CONNECTED
  There is currently no connection to the target system.
- HA_IS_NOT_OPERATIONAL
  This HyperSwap relation is not operational. The operation cannot be carried out on a non-operational HyperSwap relation.

Renaming a volume

Use the **vol_rename** command to rename a volume.

```
vol_rename vol=VolName new_name=Name
```

Parameters

Name	Type	Description	Mandatory
vol	Object name	Name of the volume to be renamed.	Y
new_name	Object name	New volume name.	Y

The new name of the volume must be unique in the system.

This command succeeds even if the new name is identical to the current name. It also succeeds regardless of the volume's lock state.

Renaming a snapshot does not change the name of its master volume. Renaming a master volume does not change the names of its associated snapshots.

If applied on a volume that is part of an IBM Hyper-Scale Mobility relation, the command has to be acknowledged by both source and destination volumes. Otherwise, a return code is returned (see below).

Example:

```
vol_rename vol=DBVolume new_name=DBVolume1
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission	Condition
Storage administrator	Allowed	N/A
Storage integration administrator	Allowed	N/A
Application administrator	Conditionally Allowed	The volume is a snapshot, where its master volume is mapped to a host or cluster associated with the user and the snapshot was created by an application administrator.
Security administrator	Disallowed	N/A
Read-only users	Disallowed	N/A
Technicians	Disallowed	N/A

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.
- VOLUME_EXISTS
  The volume name already exists.
- SNAPSHOT_IS_INTERNAL
  Internal snapshots cannot be mapped, modified or deleted.
- SNAPSHOT_IS_PART_OF_SNAPSHOT_GROUP
  The snapshot is part of a snapshot group.
- VOLUME_BAD_PREFIX
  The volume name has a reserved prefix.
- SNAPSHOT_IS_CONSISTENT_ELCS
  If a mirrored volume is not consistent, then its ELCS is protected and cannot be deleted.
- OLVM_ERROR
  IBM Hyper-Scale Mobility error.
- COMMAND_NOT_SUPPORTED_FOR_OLVM_VOLUMES
  This command is not supported for IBM Hyper-Scale Mobility volumes.
- OPERATION_DENIED_OBJECT_MANAGED
  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support

- LOCAL_PEER_IS_NOT_MASTER
  The local peer is not primary.
- TARGET_NOT_CONNECTED
  There is currently no connection to the target system.
- REMOTE_TARGET_NOT_CONNECTED
  There is currently no connection from the target system.
- HA_IS_NOT_OPERATIONAL
  This HyperSwap relation is not operational. The operation cannot be carried out on a non-operational HyperSwap relation.
- REMOTE_VOLUME_EXISTS
  The secondary volume with the indicated name already exists. The name cannot be reused.
- REMOTE_ALU_EXISTS
  An ALU with the indicated secondary volume name already exists on the remote machine.

Resizing a volume

Use the **vol_resize** command to resize a volume.

```
vol_resize vol=VolName  < size=GB | size_blocks=BLOCKS > [ shrink_volume=<yes|no> ]
 [ force_on_inactive_mirror=<yes|no> ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	The name of the volume to be resized.	Y	N/A
size	N/A	The new volume size.	N	N/A
size_blocks	N/A	New size of volumes in number of blocks.	N	N/A
shrink_volume	Boolean	Must be specified as `yes` if the new size is smaller than the current size.	N	No
force_on_inactive_mirror	Boolean	The parameter is required for a successful resize of a volume if (1) the volume is mirrored, (2) the volume is a master, and (3) the mirror has been deactivated by the system following a previously issued resize command that failed to successfully complete due to a communication error.	N	No

The volume can be resized in either direction. However, whenever the volume is downsized, you have to specify this with shrink_volume=yes.

The new size of the volume is rounded up in increments of approximately 1 GB. In some cases, rounding of up to 5% of the total volume size can take place.

If the new size equals the current size, the command succeeds without changes to the volume.

The volume's address space is extended at its end to reflect the increased size, and the additional capacity is logically formatted (that is, zeros are returned for all read commands).

When resizing a regular volume (not a writable snapshot), all storage space that is needed to support the additional volume's capacity is reserved (static allocation). This guarantees the functionality and integrity of the volume, regardless of the resource levels of the volume's storage pool. The command fails if this reservation cannot be committed.

The volume's lock state must be unlocked when the command is issued, or otherwise the command fails.

- Resizing a master volume does not change the size of its associated snapshots.
- These snapshots can still be used to restore their individual master volumes.
- A snapshot is resized in a similar way: the resize does not change the size of its master volume.

Using the force_on_inactive_mirror parameter:

- This parameter forces the resizing of a mirror peer even if mirroring is inactive (this may happen when the mirroring cannot be activated due to size mismatch).

In the following example, the *-y* option suppresses the ARE_YOU_SURE_YOU_WANT_TO_ENLARGE_VOLUME Y/N prompt.

Example:

```
vol_resize -y vol=DBVolume size=2500
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed

User Category	Permission
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_ENLARGE_VOLUME
  Are you sure you want to increase the volume size?
- ARE_YOU_SURE_YOU_WANT_TO_REDUCE_VOLUME
  Decreasing the volume size may cause data loss. Are you sure you want to proceed?
- VOLUME_SIZE_VERY_LARGE_ARE_YOU_SURE
  The volume size is very large. It may not be possible to mirror this volume to older versions of the storage system. Are you sure?

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.
- VOLUME_TOO_BIG
  No space to allocate to the volume.
- REMOTE_VOLUME_SIZE_ABOVE_LIMIT
  The specified volume size is above the limit of the remote machine.
- VOLUME_LOCKED
  The volume is locked.
- VOLUME_HAS_DATA_MIGRATION
  Data Migration is defined for this volume.
- CAN_NOT_SHRINK_MAPPED_VOLUME
  A mapped volume's size cannot be decreased.
- CAN_NOT_SHRINK_VOLUME_WITH_SNAPSHOTS
  The size of volume with snapshots cannot be decreased.
- CAN_NOT_SHRINK_REMOTE_VOLUME_WITH_SNAPSHOTS
  The remote volume has snapshots.
- CAN_NOT_SHRINK_MAPPED_REMOTE_VOLUME
  The remote volume is mapped.
- VOLUME_IS_BOUND
  The volume is bound to an ALU.

  Troubleshooting: Unbind the volume from the ALU.
- REMOTE_VOLUME_HAS_DATA_MIGRATION
  Data migration is already defined for the secondary volume.
- VOLUME_CANNOT_HAVE_ZERO_SIZE
  The volume size cannot be zero.
- CAN_NOT_SHRINK_SNAPSHOTS
  The size of snapshots cannot be decreased.
- CAN_NOT_RESIZE_ASYNC_INTERVAL_VOLUMES
  The size of volumes with asynchronous mirroring cannot be changed.
- CAN_NOT_SHRINK_VOLUME
  The size of volumes cannot be decreased without an explicit request.
- MIRROR_SIZE_MISMATCH
  The secondary and primary volume sizes are different.
- MIRROR_POSSIBLE_SIZE_MISMATCH
  The secondary and primary volume sizes may be different.
- HA_POSSIBLE_SIZE_MISMATCH
  Primary and secondary HyperSwap volume sizes may be different.
- VOLUME_SIZE_ABOVE_LIMIT
  The specified volume size is above the limit.
- COMMAND_NOT_SUPPORTED_FOR_OLVM_VOLUMES
  This command is not supported for IBM Hyper-Scale Mobility volumes.

- MIRROR_IS_NON_OPERATIONAL
  The mirror is non-operational.
- VOLUME_IS_SLAVE
  The volume is defined as a secondary volume.
- MIRROR_RETRY_OPERATION
  There is an operation in progress on this mirror.

  Troubleshooting: Retry the command in a few seconds.
- VOLUME_HAS_MULTIPLE_RELATIONS
  The volume has multiple relations. The operation is not allowed or a target must be specified.
- REMOTE_MIRROR_IS_STANDBY
  The remote mirror is marked as Standby.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support
- LOCAL_PEER_IS_NOT_MASTER
  The local peer is not primary.
- TARGET_NOT_CONNECTED
  There is currently no connection to the target system.
- REMOTE_TARGET_NOT_CONNECTED
  There is currently no connection from the target system.
- HA_IS_NOT_OPERATIONAL
  This HyperSwap relation is not operational. The operation cannot be carried out on a non-operational HyperSwap relation.
- HA_RETRY_OPERATION
  An operation is in progress on this HyperSwap relation.

  Troubleshooting: Try issuing the command again in a few seconds.
- VOLUME_CANNOT_BE_RESIZED_DUE_TO_CONTAINING_CONS_GROUP_SIZE_LIMIT
  The requested volume size will cause the containing CG to exceed the maximal CG size.

Unlocking a volume

Use the **vol_unlock** command to unlock a volume, so that it is no longer read-only and can be written to.

```
vol_unlock vol=VolName
```

Parameters

Name	Type	Description	Mandatory
vol	Object name	The name of the volume to unlock.	Y

An unlocked volume is no longer write-protected.

The lock state of regular volumes is set to *unlocked* when they are created. The lock state of snapshots is set to *locked* when they are created.

In addition to the lock state, snapshots also have a modification state. The modification state is a read-only state (which cannot be changed by the user explicitly) and it is initially set to *unmodified* when the snapshot is created. The first time a snapshot lock state is set to *unlocked*, the modification state of the snapshot is changed to *modified*, and it is never changed thereafter.

The modification time is the time when the unlock command was executed, regardless of the actual changes performed on the volume via write commands.

If applied on a volume that is part of an IBM Hyper-Scale Mobility relation, the command has to be acknowledged by both source and destination volumes. Otherwise, a return code is returned (see below).

Example:

```
vol_unlock vol=DBVolume
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission	Condition
Storage administrator	Allowed	N/A

User Category	Permission	Condition
Storage integration administrator	Allowed	N/A
Application administrator	Conditionally Allowed	The volume is a snapshot, where its master volume is mapped to a host or cluster associated with the user and the snapshot was created by an application administrator.
Security administrator	Disallowed	N/A
Read-only users	Disallowed	N/A
Technicians	Disallowed	N/A

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_UNLOCK_SNAPSHOT
  Are you sure you want to unlock snapshot *Snapshot*?

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.
- VOLUME_IS_OLVM_PROXY
  The volume is in an IBM Hyper-Scale Mobility Proxy phase.
- VOLUME_IS_SLAVE
  The volume is defined as a secondary volume.
- VOLUME_IS_OLVM_DESTINATION
  The volume is defined as an IBM Hyper-Scale Mobility destination.
- SNAPSHOT_IS_INTERNAL
  Internal snapshots cannot be mapped, modified or deleted.
- SNAPSHOT_IS_PART_OF_SNAPSHOT_GROUP
  The snapshot is part of a snapshot group.
- OPERATION_DENIED_OBJECT_MANAGED
  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support
- SYSTEM_OUT_OF_PHYSICAL_SPACE
  The operation not allowed while the system is out of physical space.
- LOCAL_PEER_IS_NOT_MASTER
  The local peer is not primary.
- TARGET_NOT_CONNECTED
  There is currently no connection to the target system.
- REMOTE_TARGET_NOT_CONNECTED
  There is currently no connection from the target system.
- HA_IS_NOT_OPERATIONAL
  This HyperSwap relation is not operational. The operation cannot be carried out on a non-operational HyperSwap relation.

Volume snapshot management commands

This section describes the command-line interface (CLI) for snapshot management.

See also:

- Volume management commands
- Consistency group management commands
- Storage pool management commands

- **Changing a snapshot deletion priority**
  Use the **snapshot_change_priority** command to change a snapshot's deletion priority.
- **Creating a snapshot**
  Use the **snapshot_create** command to create a snapshot of an existing volume.
- **Deleting a snapshot**
  Use the **snapshot_delete** command to delete a snapshot.
- **Duplicating a snapshot**
  Use the **snapshot_duplicate** command to duplicate an existing snapshot.

- **Formatting a snapshot**
  Use the **snapshot_format** command to format a snapshot.
- **Listing snapshot information**
  Use the **snapshot_list** command to list snapshot information.
- **Restoring a volume from a snapshot**
  Use the **snapshot_restore** command to restore a master volume or a snapshot from one of its associated snapshots.

Changing a snapshot deletion priority

Use the **snapshot_change_priority** command to change a snapshot's deletion priority.

```
snapshot_change_priority snapshot=SnapshotName  delete_priority=del_value
```

Parameters

Name	Type	Description	Mandatory
snapshot	Object name	Name of the snapshot whose delete_priority is to be changed.	Y
delete_priority	Integer	The priority for deleting the volume's snapshot.	Y

This command changes the priority of the deletion of an existing snapshot. The deletion priority determines which snapshots are deleted first when the system runs out of snapshot storage.

The Auto Delete Priority can have a value between 1 and 4, as follows:

- 1 = Is the last to be deleted automatically ("1" is the default set by the system)
- ...
- 4 = Is the first to be deleted automatically

Example:

```
snapshot_change_priority snapshot=DBVolume.snapshot1 delete_priority=4
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission	Condition
Storage administrator	Allowed	N/A
Storage integration administrator	Allowed	N/A
Application administrator	Conditionally Allowed	The master volume of the snapshot is mapped to a host or cluster associated with the user and the snapshot was created by the application administrator.
Security administrator	Disallowed	N/A
Read-only users	Disallowed	N/A
Technicians	Disallowed	N/A

Return codes

- VOLUME_IS_NOT_A_SNAPSHOT
  The operation is permitted on snapshots only.
- SNAPSHOT_ILLEGAL_PRIORITY
  Illegal snapshot priority; must be an integer between 1 and 4.
- SNAPSHOT_IS_INTERNAL
  Internal snapshots cannot be mapped, modified or deleted.
- SNAPSHOT_IS_PART_OF_SNAPSHOT_GROUP
  The snapshot is part of a snapshot group.
- VOLUME_BAD_NAME
  The volume name does not exist.
- SNAPSHOT_IS_CONSISTENT_ELCS
  If a mirrored volume is not consistent, then its ELCS is protected and cannot be deleted.
- OPERATION_DENIED_OBJECT_MANAGED
  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

Troubleshooting: Contact IBM Support

Creating a snapshot

Use the **snapshot_create** command to create a snapshot of an existing volume.

```
snapshot_create vol=VolName  < [ name=Name ] [ delete_priority=del_value ]
 > | < overwrite=Name > [ ext_id=Identifier ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	Name of the volumes to snapshot.	Y	N/A
name	Object name	Names of the new snapshots.	N	Auto-generated names.
delete_priority	Integer	The deletion priority of the volume's snapshot.	N	1
overwrite	Object name	Name of an existing snapshot to be overwritten with the current volumes content.	N	N/A
ext_id	String	External identifier of the volume.	N	N/A

This command creates a new snapshot for an existing volume, which is referred to as the snapshot's master volume. The snapshot's content is the same as the master volume at the exact point in time when the snapshot was created. The snapshot remains unchanged, although the master volume keeps changing after the snapshot is created. Upon a successful completion of this command, the snapshot is created and assigned a name that can later be used by other commands. The name does not have to be new. It can be the name of an already existing snapshot (in such a case, the already existing snapshot is overridden).

A write operation can be processed at the exact time of the snapshot creation, meaning that the write operation request was sent to the system before the command was executed, while the write was acknowledged after the command was executed. In this case, the content of the snapshot is not deterministic and may either contain the original value before the write operation, or the new value after the write operation. In fact, the snapshot's data may even contain a mixture of the two, where some blocks are equal to the volume before the write operation and other blocks are equal to the value after the write operation.

The new snapshot is initially locked for changes.

The created snapshot acts like a regular volume, except for the differences described below:

- The snapshot's name is either automatically generated from its master volume's name or given as a parameter to the command. It can later be changed without altering the snapshot's modification state.
- Upon successful completion of the command, the system assigns a unique SCSI ID to the snapshot. The creation time of the snapshot is set to the current time and is never changed until the snapshot is deleted.
- The size of the snapshot is the same as its master volume's size, but no storage space is reserved for the snapshot. This means that the functionality of the snapshot is not guaranteed. When the snapshot's storage pool is exhausted, the snapshot may be deleted.
- The snapshot's lock state is initially set to "locked", and as long as it is not "unlocked", the snapshot remains an exact image of the master volume at creation time and can be the source for a restore operation. The modification state of the snapshot is initially set to "unmodified".

During creation, the snapshot's deletion priority can be set explicitly, or it is automatically set to the default value. The deletion priority determines which snapshots will be deleted first when the storage pool runs out of snapshot storage. This may happen due to the redirect-on-write mechanisms which share unchanged data between volumes and their snapshots, as well as between snapshots of the same volume.

The Auto Delete Priority can have a value between 1 and 4, as follows:

- 1 = Is last to be deleted automatically ("1" is the default set by the system)
- ...
- 4 = Is first to be deleted automatically

The snapshot is associated with its master volume and this association cannot be broken or changed as long as the snapshot exists.

The **overwrite** option copies the current content of the volume into one of its existing snapshots (set as an input argument). The overwritten snapshot keeps the same SCSI device WWN and same mapping, so hosts maintain a continuous mapping to the snapshot, without any need for a rescan or similar operation. The overwritten snapshot must be an existing snapshot of the given volume. The overwritten snapshot cannot be part of a snapshot group.

This command fails when no snapshot space is defined in the storage pool the master volume belongs to.

Mirroring limitations:

- This command fails if the volume is a slave of an asynchronous mirroring coupling.
- This command fails if the volume is a slave of an inconsistent synchronous coupling.

Example:

```
snapshot_create vol=DBVolume name=DBVolume.snapshot1 delete_priority=2
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission	Condition
Storage administrator	Allowed	N/A

User Category	Permission	Condition
Storage integration administrator	Allowed	N/A
Application administrator	Conditionally Allowed	The volume is mapped to a host or a cluster associated with the user. If a snapshot overwrite is used, the target snapshot must be one created by a server administrator.
Security administrator	Disallowed	N/A
Read-only users	Disallowed	N/A
Technicians	Disallowed	N/A

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.
- MAX_VOLUMES_REACHED
  The maximum allowed number of volumes is already reached.
- DOMAIN_MAX_VOLUMES_REACHED
  The domain exceeds the maximum allowed number of volumes.
- SNAPSHOT_ILLEGAL_PRIORITY
  Illegal snapshot priority; must be an integer between 1 and 4.
- VOLUME_IS_SNAPSHOT
  The operation is not permitted on snapshots.
- VOLUME_EXISTS
  The volume name already exists.
- VOLUME_BAD_PREFIX
  The volume name has a reserved prefix.
- VOLUME_DATA_MIGRATION_UNSYNCHRONIZED
  Data Migration to this volume has not completed.
- OVERWRITE_SNAPSHOT_BAD_NAME
  The snapshot name does not exist.
- OVERWRITE_SNAPSHOT_IS_MASTER_VOL
  This snapshot cannot be overwritten because it is a primary volume.
- SNAPSHOT_OVERWRITE_MISMATCH
  The specified snapshot is not a snapshot of the specified volume.
- SNAPSHOT_IS_PART_OF_SNAPSHOT_GROUP
  The snapshot is part of a snapshot group.
- SNAPSHOT_IS_INTERNAL
  Internal snapshots cannot be mapped, modified or deleted.
- POOL_SNAPSHOT_LIMIT_REACHED
  There is not enough space to create a snapshot.
- VOLUME_IS_NOT_CONSISTENT_SLAVE
  The operation not allowed on an inconsistent secondary volume.
- VOLUME_IS_NOT_CONSISTENT_OLVM_DESTINATION
  The operation not allowed on an inconsistent IBM Hyper-Scale Mobility volume.
- VOLUME_IS_OLVM_PROXY
  The volume is in an IBM Hyper-Scale Mobility Proxy phase.
- SNAPSHOT_HAS_ACTIVE_SYNC_JOB
  The snapshot is currently the target of an active sync job.

  Troubleshooting: Please wait for the sync job to complete.
- TOO_MANY_FAST_SNAPSHOTS_IN_VOLUME
  The maximum allowed number of fast snapshots for this volume is already reached.
- OPERATION_DENIED_OBJECT_MANAGED
  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.
- NUM_VOLUMES_WILL_EXCEED_MAXIMUM
  Cannot create all the volumes, because otherwise the number of volumes will exceed the allowed maximum.
- DOMAIN_WILL_EXCEED_MAXIMUM_VOLUMES_ALLOWED
  Cannot create all the volumes, because otherwise the maximum allowed number of volumes in the domain will be exceeded.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

Troubleshooting: Contact IBM Support

- MAX_SNAPSHOTS_PER_VOLUME_REACHED
  The maximum allowed number of snapshots is already reached.

Deleting a snapshot

Use the **snapshot_delete** command to delete a snapshot.

```
snapshot_delete snapshot=SnapshotName
```

Parameters

Name	Type	Description	Mandatory
snapshot	Object name	Snapshot to be deleted.	Y

This command cannot be used to delete a master volume, or a snapshot which is mapped to a host or cluster, or an internal snapshot of a mirroring.

Example:

```
snapshot_delete snapshot=DBVolume.snapshot1
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission	Condition
Storage administrator	Allowed	N/A
Storage integration administrator	Allowed	N/A
Application administrator	Conditionally Allowed	The master volume of the snapshot is mapped to a host or cluster associated with the user and the snapshot was created by the application administrator.
Security administrator	Disallowed	N/A
Read-only users	Disallowed	N/A
Technicians	Disallowed	N/A

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.
- VOLUME_IS_NOT_A_SNAPSHOT
  The operation is permitted on snapshots only.
- SNAPSHOT_IS_INTERNAL
  Internal snapshots cannot be mapped, modified or deleted.
- SNAPSHOT_IS_PART_OF_SNAPSHOT_GROUP
  The snapshot is part of a snapshot group.
- SNAPSHOT_IS_MAPPED
  A snapshot that is mapped to a host cannot be deleted.
- VOLUME_IS_BOUND
  The volume is bound to an ALU.

  Troubleshooting: Unbind the volume from the ALU.
- SNAPSHOT_HAS_ACTIVE_SYNC_JOB
  The snapshot is currently the target of an active sync job.

  Troubleshooting: Please wait for the sync job to complete.
- SNAPSHOT_IS_CONSISTENT_ELCS
  If a mirrored volume is not consistent, then its ELCS is protected and cannot be deleted.
- OPERATION_DENIED_OBJECT_MANAGED
  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support

Duplicating a snapshot

Use the **snapshot_duplicate** command to duplicate an existing snapshot.

```
snapshot_duplicate snapshot=SnapshotName [ name=Name ]
```

Parameters

Name	Type	Description	Mandatory	Default
snapshot	Object name	The name of the snapshot to duplicate.	Y	N/A
name	Object name	Name of the new snapshot to be generated.	N	Automatically generated name.

The newly created snapshot is initially locked for changes and is associated with the master volume of the existing snapshot. The content of the newly created snapshot is identical with the content of the source snapshot.

It is useful to duplicate a snapshot before unlocking it for write operations. The duplicate snapshot can be used as a logical backup of the data in case the write operation caused logical data corruption.

Upon successful completion of the command, a new duplicate snapshot is created.

The duplicated snapshot is identical with the source snapshot. It has the same creation time and behaves as if it was created at the exact same moment and from the same master volume.

The duplicate snapshot's name is either automatically generated from its master volume's name or provided as a parameter. It can later be changed without altering its modification state.

A snapshot can be duplicated multiple times. A duplicated snapshot can be the source for further duplications.

Example:

```
snapshot_duplicate snapshot=DBVolume.snapshot1 name=DBVolume.snapshot1.copy
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission	Condition
Storage administrator	Allowed	N/A
Storage integration administrator	Allowed	N/A
Application administrator	Conditionally Allowed	The master volume of the snapshot is mapped to a host or cluster associated with the user and the snapshot was created by the application administrator.
Security administrator	Disallowed	N/A
Read-only users	Disallowed	N/A
Technicians	Disallowed	N/A

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.
- MAX_VOLUMES_REACHED
  The maximum allowed number of volumes is already reached.
- DOMAIN_MAX_VOLUMES_REACHED
  The domain exceeds the maximum allowed number of volumes.
- VOLUME_IS_NOT_A_SNAPSHOT
  The operation is permitted on snapshots only.
- VOLUME_EXISTS
  The volume name already exists.
- SNAPSHOT_IS_PART_OF_SNAPSHOT_GROUP
  The snapshot is part of a snapshot group.
- VOLUME_BAD_PREFIX
  The volume name has a reserved prefix.
- OPERATION_DENIED_OBJECT_MANAGED
  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

Troubleshooting: Contact IBM Support

- MAX_SNAPSHOTS_PER_VOLUME_REACHED
  The maximum allowed number of snapshots is already reached.

Formatting a snapshot

Use the **snapshot_format** command to format a snapshot.

```
snapshot_format snapshot=SnapshotName
```

Parameters

Name	Type	Description	Mandatory
snapshot	Object name	The snapshot to be formatted.	Y

This command deletes the content of a snapshot while maintaining its mapping to the host. The format operation results with:

- The formatted snapshot is read-only
- The format operation has no impact on performance
- The formatted snapshot does not consume space
- Reading from the formatted snapshot always returns zeroes
- The formatted snapshot can be overridden
- The formatted snapshot can be deleted
- The formatted snapshot deletion priority can be changed

Example:

```
snapshot_format snapshot
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- SNAPSHOT_IS_PART_OF_SNAPSHOT_GROUP
  The snapshot is part of a snapshot group.
- SNAPSHOT_HAS_ACTIVE_SYNC_JOB
  The snapshot is currently the target of an active sync job.

  Troubleshooting: Please wait for the sync job to complete.
- VOLUME_BAD_NAME
  The volume name does not exist.
- SNAPSHOT_IS_INTERNAL
  Internal snapshots cannot be mapped, modified or deleted.
- MAX_VOLUMES_REACHED
  The maximum allowed number of volumes is already reached.
- SNAPSHOT_IS_FORMATTED
  The snapshot is formatted.
- ELCS_CANNOT_BE_FORMATTED
  The snapshot is an ELCS and cannot be formatted.
- VOLUME_IS_NOT_A_SNAPSHOT
  The operation is permitted on snapshots only.
- OPERATION_DENIED_OBJECT_MANAGED
  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.
- DATA_REDUCTION_TIER_IS_OFFLINE

The data reduced tier is offline, the operation is not allowed.

Troubleshooting: Contact IBM Support

Listing snapshot information

Use the **snapshot_list** command to list snapshot information.

```
snapshot_list vol=VolName [ domain=DomainName ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	List of all the snapshots of this volume.	Y	N/A
domain	Object name	The domain name.	N	All Domains

This command lists snapshot information for all the snapshots of a specified volume.

It displays the following VAAI fields (available in XML output format):

- enable_VAAI
- user_disabled_VAAI

The command displays the following snapshot format field (available in XML output format):

- snapshot_format

Field ID	Field output	Default position
name	Name	1
size	Size (GB)	2
size_MiB	Size (MiB)	N/A
vol_copy_type	Copy type	3
master_name	Master Name	4
cg_name	Consistency Group	5
pool_name	Pool	6
creator	Creator	7
written	Written (GB)	8
written_MiB	Written (MiB)	N/A
proxy	Proxy	N/A
capacity	Capacity (blocks)	N/A
modified	Modified	N/A
sg_name	Snapshot Group Name	N/A
delete_priority	Deletion Priority	N/A
locked	Locked	N/A
snapshot_time	Snapshot Creation Time	N/A
snapshot_time_on_master	Master Copy Creation Time	N/A
snapshot_internal_role	Snapshot Internal Role	N/A
snapshot_of	Snapshot of	N/A
sg_snapshot_of	Snapshot of Snap Group	N/A
wwn	WWN	N/A
mirrored	Mirrored	N/A
locked_by_pool	Locked by Pool	N/A
capacity_used_by_snapshots_MiB	Capacity Used by Snapshots (MiB)	N/A
short_lived_io	Short Live IO	N/A
enable_VAAI	VAAI enabled	N/A
user_disabled_VAAI	VAAI disabled by user	N/A
snapshot_format	Snapshot Format	N/A
unmap_support	Unmap Support	N/A
managed	Managed	N/A
marked	Marked	N/A
perf_class	Performance Class Name	N/A
thin_provisioning_savings	Thin Provisioning Savings (%)	N/A
est_compression_factor	Est. Compression Factor	N/A
unique_stored_data	Unique Stored Data (GB)	N/A
ha	HA Relation	N/A
target_port_group_id	TPG ID	N/A
target_port_group_state	TPG State	N/A
lock_modes	Lock Modes	N/A
copy_master_wwn	Copy Master	N/A

Example:

```
snapshot_list vol=DBVolume
```

Output:

```
Name               Size (GB)  Master Name  Consistency Group  Pool
DBVolume.sp1       2508       DBVolume                        default
DBVolume.sp1.copy  2508       DBVolume                        default
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Restoring a volume from a snapshot

Use the **snapshot_restore** command to restore a master volume or a snapshot from one of its associated snapshots.

```
snapshot_restore snapshot=SnapshotName [ target_snapshot=SnapshotName ]
```

Parameters

Name	Type	Description	Mandatory	Default
snapshot	Object name	Name of the snapshot with which to restore its master volume, or snapshot.	Y	N/A
target_snapshot	Object name	Snapshot to be restored.	N	Restore the master volume.

This command restores the data of a master volume from one of its associated snapshots.

Issuing a restore command, logically copies the data of the source snapshot onto its volume. The volume's data is therefore restored to the state of the snapshot creation. If the volume was resized after the snapshot was created, the restore operation resizes the volume back to its original size.

All the snapshots associated with the volume are left unchanged during a restore operation.

It is possible to snapshot the volume before restoring it, so that the generated snapshot can be used and the data is not lost.

It is possible to restore another snapshot (the target snapshot) from the source snapshot. The target snapshot must be a snapshot of the same volume as the source snapshot. The target snapshot's content and size will be identical to the source snapshot's content and size. The target snapshot's lock/unlock status will remain as it was.

Restoring a mirrored volume:

- Delete the mirror
- Restore the volume
- Re-establish the mirror

It is impossible to restore a volume while it is mirrored.

Example:

```
snapshot_restore snapshot=DBVolume.snapshot1
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission	Condition
Storage administrator	Allowed	N/A
Storage integration administrator	Allowed	N/A
Application administrator	Conditionally Allowed	Both target and source are snapshots of the same master volume. This master volume is mapped to a host or cluster associated with the user, and the target snapshot was created by an application administrator.
Security administrator	Disallowed	N/A
Read-only users	Disallowed	N/A

User Category	Permission	Condition
Technicians	Disallowed	N/A

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_RESTORE_SNAPSHOT
  Are you sure you want to restore the volume from snapshot *Snapshot*?

Return codes

- VOLUME_HAS_DATA_MIGRATION
  Data Migration is defined for this volume.
- VOLUME_BAD_NAME
  The volume name does not exist.
- VOLUME_IS_NOT_A_SNAPSHOT
  The operation is permitted on snapshots only.
- VOLUME_TOO_BIG
  No space to allocate to the volume.
- SNAPSHOT_IS_PART_OF_SNAPSHOT_GROUP
  The snapshot is part of a snapshot group.
- VOLUME_HAS_MIRROR
  A mirror is defined for this volume.
- VOLUME_HAS_HA
  This operation is forbidden on a volume with a HyperSwap relation.
- VOLUME_LOCKED
  The volume is locked.
- SNAPSHOTS_BELONG_TO_DIFFERENT_MASTERS
  The target and source snapshots must be snapshots of the same volume.
- TARGET_SNAPSHOT_BAD_NAME
  The target snapshot name does not exist.
- TARGET_SNAPSHOT_IS_PART_OF_SNAPSHOT_GROUP
  The target snapshot is part of a snapshot group.
- TARGET_SNAPSHOT_IS_MASTER
  The target snapshot is a primary volume.
- TARGET_SNAPSHOT_IS_OLVM_DESTINATION
  The target snapshot is an IBM Hyper-Scale Mobility destination volume.
- TARGET_SNAPSHOT_IS_OLVM_PROXY
  The target snapshot is an IBM Hyper-Scale Mobility proxy volume.
- TARGET_SNAPSHOT_SAME_AS_SNAPSHOT
  The source snapshot must be different from the target snapshot.
- TARGET_SNAPSHOT_HAS_ACTIVE_SYNC_JOB
  The target snapshot is currently the target of an active sync job.

  Troubleshooting: Please wait for sync job to complete
- OPERATION_DENIED_OBJECT_MANAGED
  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support

Consistency group management commands

This section describes the command-line interface (CLI) for consistency group management.

See also:

- Volume management commands
- Volume snapshot management commands
- Storage pool management commands

- **Adding a volume to a consistency group**
  Use the **cg_add_vol** command to add a volume to a consistency group.
- **Creating consistency groups**
  Use the **cg_create** command to create a consistency group.
- **Deleting a consistency group**
  Use the **cg_delete** command to delete a consistency group.
- **Listing consistency groups**
  Use the **cg_list** command to list consistency groups.
- **Removing a volume from a consistency group**
  Use the command **cg_remove_vol** to remove a volume from a consistency group.
- **Renaming a consistency group**
  Use the **cg_rename** command to rename consistency groups.
- **Creating a cross-system consistency group**
  Use the **xcg_create** command to create a cross-system consistency group (XCG) definition.
- **Associating an existing consistency group with a cross-system consistency group definition**
  Use the **xcg_add_cg** command to associate an existing consistency group to a cross-system consistency group definition.
- **Removing an existing consistency group from a cross-system consistency group definition**
  Use the **xcg_remove_cg** command to remove an existing consistency group from a cross-system consistency group definition.
- **Adding a remote system name to a cross-system consistency group definition**
  Use the **xcg_add_remote_system** command to add a remote system name to a cross-system consistency group definition.
- **Removing a remote system name from a cross-system consistency group definition**
  Use the **xcg_remove_remote_system** command to remove a remote system name from a cross-system consistency group definition.
- **Listing cross-system consistency group definitions**
  Use the **xcg_get_local_cgs** command to list cross-system consistency group definitions together with the contained consistency groups.
- **Retrieving remote systems in a specified cross-system consistency group**
  Use the **xcg_get_remote_systems** command to retrieve the names of remote systems that are a part of the specified cross-system consistency group.
- **Deleting a cross-system consistency group**
  Use the **xcg_delete** command to delete a cross-system consistency group (XCG) definition.
- **Listing cross-system consistency group definitions**
  Use the **xcg_list** command to list cross-system consistency group definitions together with the contained consistency groups.

Adding a volume to a consistency group

Use the **cg_add_vol** command to add a volume to a consistency group.

```
cg_add_vol cg=cgName vol=VolName
```

Parameters

Name	Type	Description	Mandatory
cg	Object name	Name of a consistency group.	Y
vol	Object name	Name of the volume to be added.	Y

The consistency group can contain up to 128 volumes. The maximum allowed total effective capacity of the volumes in a consistency group is limited to 1 PB.

Requirements for successful command completion:

- The volume and consistency group are associated with the same pool.
- The volume is not already part of a consistency group.
- The volume is not a snapshot.
- The consistency group has less than the maximum number of volumes (see above).

Adding a mirrored volume to a non-mirrored consistency group:

- Such an addition always succeeds and the volume will retain its mirroring settings.

Requirements for successful command completion for a mirrored consistency group:

- The command must be issued only on the master consistency group.
- The command cannot be run during the initialization of the volume or consistency group.
- The volume does not have any outstanding ad-hoc sync jobs.
- The volume has to be mirrored, and its following mirroring settings must be identical to those of the consistency group: mirroring type (for example, synchronous), mirroring status, mirroring target, target pool, designation.
- In addition, for a mirrored consistency group that is defined as `sync_best_effort` (synchronous):
  - The synchronization status of both volume and consistency group has to be Synchronized.
- For a mirrored consistency group that is defined as `async_interval` (asynchronous):
  - The volume and consistency group must have the following identical settings and values: schedule, remote schedule, timestamp of the last replicated snapshot.
  - The synchronization status of the volume and consistency group must be `RPO_OK`
- The link has to be up.

Adding a mirrored volume to a mirrored volume and consistency group also adds the volume's peer to the volume and consistency group's peer. Once added, the mirrored volume will be set the RPO of the mirrored volume and consistency group.

The mirrored consistency group has one sync job for all pertinent mirrored volumes within the consistency group.

If the command cg_add_vol is issued on a mirrored master consistency group, which fails to receive an acknowledgment from the slave until the command times out or due to an unexpected failure, the MIRROR_POSSIBLE_CONS_GROUP_MEMBERSHIP_MISMATCH completion code is returned. The completion code indicates that the member lists of the mirror consistency group peers might not be the same.

If the volume to be added is part of an IBM Hyper-Scale Mobility relation, the command is applicable only to the destination volume. For the command to complete successfully, the destination volume must be in a proxy state.

Note:
In rare cases, when adding a volume to a consistency group (CG) with the cg_add_vol command, the operation might fail with the return code MIRROR_LAST_SYNC_TIMES_DIFFER recorded. To avoid this issue, retry the cg_add_vol command after the next volume mirror sync job completes, by which time the synchronization times should match.

Example:

```
cg_add_vol cg=DBGroup vol=DBLog
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- CONS_GROUP_MIRROR_DESIGNATION_MISMATCH
  The volume's role in a mirroring or HyperSwap relation is different from the consistency group role. Are you sure you want to add the volume to that consistency group?

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.
- CONS_GROUP_BAD_NAME
  The consistency group name does not exist.
- CONS_GROUP_IS_SLAVE
  The consistency group's role in a mirroring relationship is secondary.
- CONS_GROUP_MIRROR_STATE_MISMATCH
  All volumes in a mirrored consistency group must have the same mirroring sync state.
- MAX_VOLUMES_IN_CONS_GROUP_REACHED
  The consistency group contains the maximum allowed number of volumes.
- MAX_VOLUMES_IN_REMOTE_CONS_GROUP_REACHED
  The remote consistency group contains the maximum allowed number of volumes.
- MIRROR_CONS_GROUP_CANNOT_CONTAIN_NON_MIRRORED_VOLS
  A mirrored consistency group cannot include volumes that are not mirrored.
- MIRROR_HAS_SYNC_JOB
  The operation is not permitted on a mirror with active sync jobs.
- MIRROR_IS_NOT_SYNCHRONIZED
  The mirror is not synchronized.
- MIRROR_LAST_SYNC_TIMES_DIFFER
  All mirrors must have the same last sync time.
- MIRROR_RETRY_OPERATION
  There is an operation in progress on this mirror.

  Troubleshooting: Retry the command in a few seconds.
- REMOTE_VOLUME_BAD_POOL
  The remote volume and remote consistency group belong to different storage pools.
- REMOTE_VOLUME_BELONGS_TO_CONS_GROUP
  The remote volume belongs to a consistency group.

- TARGET_NOT_CONNECTED
  There is currently no connection to the target system.
- VOLUME_BAD_POOL
  The volume belongs to a different storage pool.
- VOLUME_BELONGS_TO_CG
  The volume belongs to a consistency group.
- VOLUME_DATA_MIGRATION_UNSYNCHRONIZED
  Data Migration to this volume has not completed.
- VOLUME_IS_SNAPSHOT
  The operation is not permitted on snapshots.
- CONS_GROUP_MIRROR_SCHEDULE_MISMATCH
  All volumes in a mirrored consistency group must have the same mirroring schedule.
- CONS_GROUP_MIRROR_TARGET_MISMATCH
  All volumes in a mirrored consistency group must have the same mirroring target.
- CONS_GROUP_MIRROR_ROLE_MISMATCH
  All volumes in a mirrored consistency group must have the same mirroring role.
- CONS_GROUP_MIRROR_ACTIVATION_MISMATCH
  All volumes in a mirrored consistency group must have the same mirroring activation state.
- CONS_GROUP_MIRROR_STANDBY_MISMATCH
  All volumes in a mirrored consistency group must have the same mirroring standby state.
- HA_HIGH_AVAILABILITY_DISABLED_IN_VOL
  The consistency group's high availability is enabled, but the volume's high availability is disabled.
- HA_HIGH_AVAILABILITY_ENABLED_IN_VOL
  The consistency group's high availability is disabled but the volume's high availability is enabled.
- CONS_GROUP_HA_ROLE_MISMATCH
  All volumes in a HyperSwap consistency group must have the same mirroring role.
- HA_LAST_SYNC_TIMES_DIFFER
  All HyperSwap relations in a consistency group must have the same last sync time.
- HA_POSSIBLE_CONS_GROUP_MEMBERSHIP_MISMATCH
  The HyperSwap consistency group may contain different volumes on the primary and secondary machines.
- CONS_GROUP_HA_ACTIVATION_MISMATCH
  All volumes in a HyperSwap consistency group must have the same HyperSwap activation state.
- CONS_GROUP_HA_TARGET_MISMATCH
  All volumes in a mirrored consistency group must have the same HyperSwap target.
- HA_RETRY_OPERATION
  An operation is in progress on this HyperSwap relation.

  Troubleshooting: Try issuing the command again in a few seconds.
- HA_IS_NOT_SYNCHRONIZED
  The HyperSwap relation is not synchronized.
- REMOTE_CONS_GROUP_MIRROR_SCHEDULE_MISMATCH
  All volumes in a mirrored consistency group on the remote machine must have identical mirroring schedule.
- CONS_GROUP_MIRROR_TYPE_MISMATCH
  All volumes in a mirrored consistency group must be of the same mirroring type.
- MIRROR_POSSIBLE_CONS_GROUP_MEMBERSHIP_MISMATCH
  The mirrored consistency group contains different volumes on the primary and secondary machines. This problem occurs whenever the cg_add_vol command was issued, but the primary machine did not receive an acknowledgment from the secondary machine until the command timed out, or due to any other unexpected failure.
- REMOTE_CONS_GROUP_CRASH_CONSISTENCY_MISMATCH
  Crash consistency of the volume does not match the state of other volumes in the group on the remote machine.
- CONS_GROUP_CRASH_CONSISTENCY_MISMATCH
  Crash consistency of the volume does not match the state of other volumes in the group.
- OPERATION_DENIED_OBJECT_MANAGED
  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.
- VOLUME_HAS_OLVM
  An IBM Hyper-Scale Mobility relationship is defined for this volume.
- VOLUME_HAS_MULTIPLE_RELATIONS

The volume has multiple relations. The operation is not allowed or a target must be specified.

- REMOTE_MIRROR_IS_STANDBY
  The remote mirror is marked as Standby.
- VOLUME_CANNOT_BE_ADDED_TO_CONS_GROUP_DUE_TO_CONS_GROUP_SIZE_LIMIT
  Addition of the volume to the CG will cause it to exceed the maximal CG size.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support
- MULTISITE_CONS_GROUP_CANNOT_CONTAIN_NON_MULTISITEED_VOLS
  A Multi-site consistency group cannot include volumes that are not in a Multi-site relation.
- MULTISITE_CONS_GROUP_ROLE_IS_NOT_MASTER
  The local consistency group is not primary in the Multi-site relation.
- MULTISITE_IS_NOT_OPERATIONAL
  A Multi-site must be operational.
- MULTISITE_IS_PARTIALY_DEFINED
  The Multi-site relation is partially defined.
- MULTISITE_CONS_GROUP_STANDBY_CONFIGURATION_MISMATCH
  The standby state of the Multi-site consistency group does not match the standby configuration of a Multi-site volume.
- MIRROR_OR_HYPERSWAP_CONS_GROUP_CANNOT_CONTAIN_MULTISITEED_VOLS
  A mirroring or HyperSwap consistency group cannot include volumes that are in a Multi-site relation.
- MULTISITE_ROLE_IS_STANDALONE_MASTER
  Command failed because a Multi-site STANDALONE master cannot be activated.
- MULTISITE_ROLE_IS_NOT_MASTER
  The local peer is not the Master in the Multi-site relation.
- MULTISITE_CONS_GROUP_LAST_SMASTER_SYNC_MISMATCH
  The state of the Multi-site consistency group does not match the last SMaster sync of a Multi-site volume.
- MULTISITE_CG_ADD_VOL_FAILED_TO_ROLLBACK_MANUAL_CLEANUP_REQUIRED
  An error occured while adding a volume to the Multi-site consistency group, and the system failed to roll back. To recover, manually ensure that consistency group constituents are identical on all systems.

Creating consistency groups

Use the **cg_create** command to create a consistency group.

```
cg_create cg=cgName pool=PoolName
```

Parameters

Name	Type	Description	Mandatory
cg	Object name	Name of the consistency group.	Y
pool	Object name	Storage pool of the consistency group.	Y

This command creates a consistency group. A consistency group is a group of volumes that can all be snapshotted at the same point of time. This is essential for snapshotting several volumes used by the same application or by applications that interact with each other in order to generate a consistent set of snapshots.

The name of the consistency group must be unique in the system.

The storage pool of the consistency group must be specified.

The consistency group is initially empty, containing no volumes.

A consistency group always belongs to a specific storage pool. All the volumes in the consistency group belong to the same storage pool as the consistency group itself.

The consistency group can be mirrored as a whole (see Creating a mirroring definition).

Example:

```
cg_create pool=p_1 cg=DBgroup
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- CONS_GROUP_NAME_EXISTS
  The consistency group name already exists.
- MAX_CONS_GROUPS_REACHED
  The maximum allowed number of consistency groups is already reached.
- POOL_DOES_NOT_EXIST
  The storage pool does not exist.
- DOMAIN_MAX_CONS_GROUPS_REACHED
  The domain exceeds the maximum allowed number of consistency groups.
- OPERATION_DENIED_OBJECT_MANAGED
  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support

Deleting a consistency group

Use the **cg_delete** command to delete a consistency group.

```
cg_delete cg=cgName
```

Parameters

Name	Type	Description	Mandatory
cg	Object name	Name of the consistency group to be deleted.	Y

This command fails if:

- The consistency group is not empty, that is, it still contains volumes.
- The consistency group is mirrored, even if it is empty.

All snapshot groups associated with the consistency group are disbanded. That means that the snapshots contained in these snapshot groups become independent snapshots.

Example:

```
cg_delete cg=DBvolumes
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- CONS_GROUP_BAD_NAME
  The consistency group name does not exist.
- CONS_GROUP_NOT_EMPTY

This operation is only allowed on an empty consistency group.

- CONS_GROUP_HAS_MIRROR
  Mirroring is defined for this consistency group.
- CONS_GROUP_BELONGS_TO_XCG
  The consistency group belongs to another cross-system consistency group.
- OPERATION_DENIED_OBJECT_MANAGED
  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support

Listing consistency groups

Use the **cg_list** command to list consistency groups.

```
cg_list [ cg=cgName ] [ managed=<yes|no|all> ] [ domain=DomainName ]
```

Parameters

Name	Type	Description	Mandatory	Default
cg	Object name	Name of a consistency group.	N	All
managed	Boolean	Determines whether to show unmanaged consistency groups (no), managed consistency groups (yes) or both (all).	N	no
domain	Object name	The domain name.	N	All Domains

This command lists the specified details for all consistency groups. If a consistency group name is indicated, only this consistency group is listed.

The listed details include the following fields:

- Name
- Mirrored - Indicates whether the consistency group is mirrored
  - Available values: Yes|No
- GP Based - Indicates whether the consistency group is based on a grouped pool
  - Values - Yes|No
- Mirror sync status - Indicates the mirroring status
  - Available values: RPO_OK|RPO_Lagging
- CG role - Indicates the peer's role
  - Available values: master|slave

Field ID	Field output	Default position
name	Name	1
pool	Pool Name	2
mirrored	Mirrored	N/A
ha	HA	N/A
managed	Managed	N/A

Example:

```
cg_list cg=DBgroup
```

Output:

```
Name     Pool Name Mirrored GP Based
DBgroup  default   Yes      No
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Removing a volume from a consistency group

Use the command **cg_remove_vol** to remove a volume from a consistency group.

```
cg_remove_vol vol=VolName
```

Parameters

Name	Type	Description	Mandatory
vol	Object name	Name of the volume to be removed.	Y

This command removes a volume from a consistency group.

A consistency group's name is deduced from the volume name. A unique name is ensured because each volume belongs to only a single consistency group. Future snapshot groups created from this consistency group will not include the snapshot associated with the removed volume.

All the snapshots of the removed volume that were created as part of this consistency group will be permanently removed from the snapshot groups they were associated with.

Following the volume removal:

- The corresponding peer volume is removed from the peer consistency group. If the consistency group is mirrored, the mirroring definition of the removed volume is retained (based on the same settings as the consistency group from which it was removed).
- The peer volume is also removed from the peer consistency group.
- The removed mirrored volume acquires the RPO of the mirrored consistency group from which it was removed.
- An event is generated.

This command succeeds even if the volume is not included in any consistency group.

Requirements for a successful command completion:

- The command can be issued only on the master.
- The link has to be up.
- The consistency group cannot have ongoing sync jobs.

If the command is issued on a mirrored contingency group master, and the master does not receive an acknowledgment from the slave because the command times out or due to an unexpected failure, a return code is returned: (MIRROR_POSSIBLE_CONS_GROUP_MEMBERSHIP_MISMATCH).

If applied on a volume that is part of an IBM Hyper-Scale Mobility relation:

- The command is applicable only to a destination volume.
- This destination volume has to be in Proxy state.

Otherwise, a completion code is returned (see below).
Example:

```
cg_remove_vol vol=DBLog
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_REMOVE_VOLUME_FROM_CONS_GROUP
  Are you sure you want to remove volume '*Volume*' from its consistency group?

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.
- VOLUME_NOT_IN_CONS_GROUP
  The volume does not belong to a consistency group.
- TARGET_NOT_CONNECTED
  There is currently no connection to the target system.
- VOLUME_IS_SNAPSHOT
  The operation is not permitted on snapshots.

- CONS_GROUP_IS_SLAVE
  The consistency group's role in a mirroring relationship is secondary.
- MIRROR_RETRY_OPERATION
  There is an operation in progress on this mirror.

  Troubleshooting: Retry the command in a few seconds.
- MIRROR_POSSIBLE_CONS_GROUP_MEMBERSHIP_MISMATCH
  The mirrored consistency group contains different volumes on the primary and secondary machines. This problem occurs whenever the cg_add_vol command was issued, but the primary machine did not receive an acknowledgment from the secondary machine until the command timed out, or due to any other unexpected failure.
- HA_POSSIBLE_CONS_GROUP_MEMBERSHIP_MISMATCH
  The HyperSwap consistency group may contain different volumes on the primary and secondary machines.
- VOLUME_IS_NOT_CONSISTENT_SLAVE
  The operation not allowed on an inconsistent secondary volume.
- SNAPSHOT_HAS_ACTIVE_SYNC_JOB
  The snapshot is currently the target of an active sync job.

  Troubleshooting: Please wait for the sync job to complete.
- OPERATION_DENIED_OBJECT_MANAGED
  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.
- VOLUME_HAS_OLVM
  An IBM Hyper-Scale Mobility relationship is defined for this volume.
- REMOTE_MIRROR_IS_STANDBY
  The remote mirror is marked as Standby.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support
- MAX_SNAPSHOTS_PER_VOLUME_REACHED
  The maximum allowed number of snapshots is already reached.
- HA_RETRY_OPERATION
  An operation is in progress on this HyperSwap relation.

  Troubleshooting: Try issuing the command again in a few seconds.
- MIRROR_HAS_SYNC_JOB
  The operation is not permitted on a mirror with active sync jobs.

Renaming a consistency group

Use the **cg_rename** command to rename consistency groups.

```
cg_rename cg=cgName new_name=Name
```

Parameters

Name	Type	Description	Mandatory
cg	Object name	The name of the consistency group to be renamed.	Y
new_name	Object name	The new name of the consistency group.	Y

The new name of the consistency group must be unique in the system.

This command succeeds even if the new name is identical with the current name.

Example:

```
cg_rename cg=DBgroup new_name=DBvolumes
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed

User Category	Permission
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- CONS_GROUP_BAD_NAME
  The consistency group name does not exist.
- CONS_GROUP_NAME_EXISTS
  The consistency group name already exists.
- OPERATION_DENIED_OBJECT_MANAGED
  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support

Creating a cross-system consistency group

Use the **xcg_create** command to create a cross-system consistency group (XCG) definition.

```
xcg_create xcg=XcgName
```

Parameters

Name	Type	Description	Mandatory
xcg	Object name	The name of the new cross-system consistency group.	Y

This command creates a cross-system consistency group (XCG) definition, with which consistency groups on different systems can be associated.

Example:

```
xcg_create xcg=DBbackup
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission	Condition
Storage administrator	Allowed	N/A
Storage integration administrator	Allowed	N/A
Application administrator	Conditionally Allowed	At least one of the volumes in the group is mapped to a host or cluster associated with the user. If a Snapshot Group overwrite is used, then the target Snapshot Group must be one created by a server administrator.
Security administrator	Disallowed	N/A
Read-only users	Disallowed	N/A
Technicians	Disallowed	N/A

Return codes

- XCG_NAME_EXISTS
  The cross-system consistency group name already exists.
- MAX_XCGS_REACHED
  The maximum allowed number of cross-system consistency groups is already reached.

Associating an existing consistency group with a cross-system consistency group definition

Use the **xcg_add_cg** command to associate an existing consistency group to a cross-system consistency group definition.

```
xcg_add_cg xcg=XcgName cg=cgName
```

Parameters

Name	Type	Description	Mandatory
xcg	Object name	Name of a cross-system consistency group.	Y
cg	Object name	Name of a consistency group.	Y

Example:

```
xcg_add_cg xcg=DBbackup cg=CGbackup
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission	Condition
Storage administrator	Allowed	N/A
Storage integration administrator	Allowed	N/A
Application administrator	Conditionally Allowed	At least one of the volumes in the group is mapped to a host or cluster associated with the user. If a Snapshot Group overwrite is used, then the target Snapshot Group must be one created by a server administrator.
Security administrator	Disallowed	N/A
Read-only users	Disallowed	N/A
Technicians	Disallowed	N/A

Return codes

- XCG_BAD_NAME
  The cross-system consistency group name does not exist.
- MAX_CONS_GROUPS_IN_XCG_REACHED
  The cross-system consistency group contains the maximum allowed number of consistency groups.
- CONS_GROUP_IS_SLAVE
  The consistency group's role in a mirroring relationship is secondary.
- CONS_GROUP_BAD_NAME
  The consistency group name does not exist.
- CONS_GROUP_ALREADY_IN_XCG
  The consistency group already belongs to a cross-system consistency group.
- CONS_GROUP_BELONGS_TO_XCG
  The consistency group belongs to another cross-system consistency group.

Removing an existing consistency group from a cross-system consistency group definition

Use the **xcg_remove_cg** command to remove an existing consistency group from a cross-system consistency group definition.

```
xcg_remove_cg xcg=XcgName cg=cgName
```

Parameters

Name	Type	Description	Mandatory
xcg	Object name	Name of a Cross-system Consistency Group.	Y
cg	Object name	Name of a Consistency Group.	Y

Example:

```
xcg_remove_cg xcg=DBbackup cg=CGBackup
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission	Condition
Storage administrator	Allowed	N/A
Storage integration administrator	Allowed	N/A
Application administrator	Conditionally Allowed	At least one of the volumes in the group is mapped to a host or cluster associated with the user. If a Snapshot Group overwrite is used, then the target Snapshot Group must be one created by a server administrator.
Security administrator	Disallowed	N/A
Read-only users	Disallowed	N/A
Technicians	Disallowed	N/A

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_REMOVE_CONS_GROUP_FROM_XCG
  Are you sure you want to remove consistency group '*CG*' from its cross-system consistency group?

Return codes

- XCG_BAD_NAME
  The cross-system consistency group name does not exist.
- CONS_GROUP_BAD_NAME
  The consistency group name does not exist.
- XCG_IS_EMPTY
  The consistency group is empty.
- CONS_GROUP_NOT_IN_XCG
  The consistency group does not belong to a cross-system consistency group.

Adding a remote system name to a cross-system consistency group definition

Use the **xcg_add_remote_system** command to add a remote system name to a cross-system consistency group definition.

```
xcg_add_remote_system xcg=XcgName remote_system=RemoteSystem
```

Parameters

Name	Type	Description	Mandatory
xcg	Object name	Name of a cross-system consistency group.	Y
remote_system	String	Name of a remote system.	Y

Example:

```
xcg_add_remote_system xcg=DBbackup remote_system=CGbackup
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission	Condition
Storage administrator	Allowed	N/A
Storage integration administrator	Allowed	N/A
Application administrator	Conditionally Allowed	At least one of the volumes in the group is mapped to a host or cluster associated with the user. If a Snapshot Group overwrite is used, then the target Snapshot Group must be one created by a server administrator.
Security administrator	Disallowed	N/A
Read-only users	Disallowed	N/A
Technicians	Disallowed	N/A

Return codes

- XCG_BAD_NAME
  The cross-system consistency group name does not exist.
- MAX_REMOTE_SYSTEMS_IN_XCG_REACHED
  The cross-system consistency group contains the maximum number of remote systems.
- REMOTE_SYSTEM_ALREADY_ADDED
  The remote system belongs to a cross-system consistency group.

Removing a remote system name from a cross-system consistency group definition

Use the **xcg_remove_remote_system** command to remove a remote system name from a cross-system consistency group definition.

```
xcg_remove_remote_system xcg=XcgName remote_system=RemoteSystem
```

Parameters

Name	Type	Description	Mandatory
xcg	Object name	Name of a Cross-system Consistency Group.	Y
remote_system	String	Name of a remote system.	Y

Example:

```
xcg_remove_remote_system xcg=DBbackup remote_system=CGbackup
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission	Condition
Storage administrator	Allowed	N/A
Storage integration administrator	Allowed	N/A
Application administrator	Conditionally Allowed	At least one of the volumes in the group is mapped to a host or cluster associated with the user. If a Snapshot Group overwrite is used, then the target Snapshot Group must be one created by a server administrator.
Security administrator	Disallowed	N/A
Read-only users	Disallowed	N/A
Technicians	Disallowed	N/A

Return codes

- XCG_BAD_NAME
  The cross-system consistency group name does not exist.
- REMOTE_SYSTEM_NOT_IN_XCG
  The remote system does not belong to a cross-system consistency group.

Listing cross-system consistency group definitions

Use the **xcg_get_local_cgs** command to list cross-system consistency group definitions together with the contained consistency groups.

```
xcg_get_local_cgs [ xcg=XcgName ]
```

Parameters

Name	Type	Description	Mandatory	Default
xcg	Object name	Name of a cross-system consistency group.	N	All Cross-system Consistency Groups.

Example:

```
xcg_get_local_cgs
```

Output:

```
Command completed successfully.
```

Field ID	Field output	Default position
name	Name	1
xcg	XCG Name	2

Access control

User Category	Permission	Condition
Storage administrator	Allowed	N/A
Storage integration administrator	Allowed	N/A
Application administrator	Conditionally Allowed	At least one of the volumes in the group is mapped to a host or cluster associated with the user. If a Snapshot Group overwrite is used, then the target Snapshot Group must be one created by a server administrator.
Security administrator	Disallowed	N/A
Read-only users	Allowed	N/A
Technicians	Disallowed	N/A

Return codes

- XCG_BAD_NAME
  The cross-system consistency group name does not exist.

Retrieving remote systems in a specified cross-system consistency group

Use the **xcg_get_remote_systems** command to retrieve the names of remote systems that are a part of the specified cross-system consistency group.

```
xcg_get_remote_systems xcg=XcgName
```

Parameters

Name	Type	Description	Mandatory
xcg	Object name	Name of a Cross-system Consistency Group.	Y

Example:

```
xcg_get_remote_systems xcg=blabla
```

Output:

```
Command completed successfully.
```

Field ID	Field output	Default position
name	Name	1
xcg	XCG Name	2

Access control

User Category	Permission	Condition
Storage administrator	Allowed	N/A
Storage integration administrator	Allowed	N/A
Application administrator	Conditionally Allowed	At least one of the volumes in the group is mapped to a host or cluster associated with the user. If a Snapshot Group overwrite is used, then the target Snapshot Group must be one created by a server administrator.
Security administrator	Disallowed	N/A
Read-only users	Allowed	N/A
Technicians	Disallowed	N/A

Return codes

- XCG_BAD_NAME
  The cross-system consistency group name does not exist.

Deleting a cross-system consistency group

Use the **xcg_delete** command to delete a cross-system consistency group (XCG) definition.

```
xcg_delete xcg=XcgName
```

Parameters

Name	Type	Description	Mandatory
xcg	Object name	Name of a cross-system consistency group.	Y

Example:

```
xcg_delete xcg=DBbackup
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission	Condition
Storage administrator	Allowed	N/A
Storage integration administrator	Allowed	N/A
Application administrator	Conditionally Allowed	At least one of the volumes in the group is mapped to a host or cluster associated with the user. If a Snapshot Group overwrite is used, then the target Snapshot Group must be one created by a server administrator.
Security administrator	Disallowed	N/A
Read-only users	Disallowed	N/A
Technicians	Disallowed	N/A

Return codes

- XCG_BAD_NAME
  The cross-system consistency group name does not exist.
- XCG_NOT_EMPTY
  The consistency group is not empty.

Listing cross-system consistency group definitions

Use the **xcg_list** command to list cross-system consistency group definitions together with the contained consistency groups.

```
xcg_list [ xcg=XcgName ]
```

Parameters

Name	Type	Description	Mandatory	Default
xcg	Object name	Name of a Cross-system Consistency Group.	N	All Cross-system Consistency Groups.

Field ID	Field output	Default position
name	Name	1
num_of_cgs	Num Of CGs	2
num_of_remote_systems	Num Of Remote Systems	3

Example:

```
xcg_list
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission	Condition
Storage administrator	Allowed	N/A

User Category	Permission	Condition
Storage integration administrator	Allowed	N/A
Application administrator	Conditionally Allowed	At least one of the volumes in the group is mapped to a host or cluster associated with the user. If a Snapshot Group overwrite is used, then the target Snapshot Group must be one created by a server administrator.
Security administrator	Disallowed	N/A
Read-only users	Allowed	N/A
Technicians	Disallowed	N/A

Snapshot set management commands

This section describes the command-line interface (CLI) for snapshot set management.

See also:

- Volume management commands
- Volume snapshot management commands
- Consistency group management commands

- **Snapshotting a consistency group**
  Use the **cg_snapshots_create** command to create a snapshot group of a consistency group.
- **Changing a snapshot group deletion priority**
  Use the **snap_group_change_priority** command to change the deletion priority of a snapshot group.
- **Deleting a snapshot group**
  Use the **snap_group_delete** command to delete a snapshot group and all its snapshots.
- **Disbanding a snapshot group**
  Use the **snap_group_disband** command to disband a snapshot group into independent snapshots.
- **Duplicating a snapshot group**
  Use the **snap_group_duplicate** command to duplicate an existing snapshot group.
- **Formatting a snapshot group**
  Use the **snap_group_format** command to format a snapshot group.
- **Listing snapshot groups**
  Use the **snap_group_list** command to list all snapshot groups or a specific one.
- **Locking a snapshot group**
  Use the **snap_group_lock** command to lock a snapshot group by locking all its snapshots.
- **Renaming a snapshot group**
  Use the **snap_group_rename** command to rename a snapshot group.
- **Restoring a consistency group from a snapshot group**
  Use the **snap_group_restore** command to restore the master volumes of a consistency group, or of a snapshot group from one of its associated snapshot groups.
- **Unlocking a snapshot group**
  Use the **snap_group_unlock** command to unlock a snapshot group by unlocking all its snapshots.
- **Setting a snapshot group descriptor**
  Use the **snap_group_set_descriptor** command to set a snapshot group descriptor.
- **Returning a snapshot group's descriptor**
  Use the **snap_group_get_descriptor** command to return a snapshot group's descriptor.

Snapshotting a consistency group

Use the **cg_snapshots_create** command to create a snapshot group of a consistency group.

```
cg_snapshots_create
cg=cgName < [ snap_group=SnapshotGroupName ]
[ delete_priority=del_value ]
 [ auto_resume=token_id ] > | <overwrite=Name>
```

Parameters

Name	Type	Description	Mandatory	Default
cg	Object name	The name of the consistency group whose snapshot will be created.	Y	N/A
snap_group	Object name	The name of the newly created snapshot group.	N	Automatically generated name.
delete_priority	Integer	The priority for deleting this volume when the system runs out of snapshot space.	N	1
overwrite	Object name	An existing snapshot group that will be overwritten with the current content.	N	N/A

Name	Type	Description	Mandatory	Default
auto_resume	Positive integer	Defines whether to resume IO to the consistency group by providing the token ID.	N	0

This command creates a consistent snapshot group of a consistency group. The snapshot group includes a snapshot for each of the volumes contained in the consistency group.

Logically, this command is comprised of the following steps:

- Suspending all I/O activity on all the volumes in the group and waiting for all pending I/Os to complete.
- Creating a snapshot for each volume in the group.
- Resuming I/O activity on all the volumes.

The main advantage of using this command (as opposed to a manual procedure) is that all snapshots are taken at the same point of time, thus ensuring that they are consistent with each other.

The snapshots in the created snapshot group are consistent with each other in the following aspects:

- They are created synchronously at the same point of time.
- All I/Os to the consistency group's volumes that were completed prior to this point of time are recorded in the snapshot's image.
- Neither I/O that was completed after this point of time is recorded in the snapshot's image.

In addition to their regular attributes, all the snapshots in the snapshot group are also associated with the consistency group.

The name of the snapshot group is either automatically generated or provided in the command line.

The delete priority of the snapshots in the snapshot group can also be provided (see Creating a snapshot). The delete priority controls which snapshots or snapshot groups are deleted first when the system runs out of space for snapshots.

The **`overwrite`** option causes the current content of the consistency group to be copied into one of its existing snapshot groups (indicated as parameter's argument). The snapshots of the overwritten snapshot group keep the same SCSI device WWN and same mapping, so hosts maintain a continuous mapping of the snapshots, and a rescan or similar operation is not needed. The overwritten snapshot group must be an existing snapshot group of the respective consistency group.

This command fails if no snapshot space is defined for the storage pool containing the consistency group.

This command fails if one or more of the volumes in the consistency group are slaves in the synchronous mirroring, and the synchronous mirroring is currently inconsistent due to either a re-synchronization or an initialization process.

Mirroring limitations:

- This command fails if the volume is a slave of an asynchronous mirroring coupling (either synchronous or asynchronous).
- This command fails if the volume is a slave of an inconsistent synchronous coupling.

Example:

```
cg_snapshots_create cg=DBgroup snap_group=DBbackupdaily
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission	Condition
Storage administrator	Allowed	N/A
Storage integration administrator	Allowed	N/A
Application administrator	Conditionally Allowed	At least one of the volumes in the group is mapped to a host or cluster associated with the user. If a Snapshot Group overwrite is used, then the target Snapshot Group must be one created by a server administrator.
Security administrator	Disallowed	N/A
Read-only users	Disallowed	N/A
Technicians	Disallowed	N/A

Return codes

- MAX_VOLUMES_REACHED
  The maximum allowed number of volumes is already reached.
- DOMAIN_MAX_VOLUMES_REACHED
  The domain exceeds the maximum allowed number of volumes.
- CONS_GROUP_BAD_NAME
  The consistency group name does not exist.
- SNAPSHOT_GROUP_BAD_NAME
  The snapshot group name does not exist.

- SNAPSHOT_GROUP_BAD_PREFIX
  The snapshot group name has a reserved prefix.
- SNAPSHOT_GROUP_NAME_EXISTS
  The snapshot group name already exists.
- CONS_GROUP_EMPTY
  The operation is not allowed on an empty consistency group.
- CONS_GROUP_MISMATCH
  The snapshot group does not match the consistency group volumes.
- OVERWRITE_SNAPSHOT_GROUP_DOES_NOT_BELONG_TO_GIVEN_GROUP
  The snapshot group belongs to another consistency group.
- POOL_SNAPSHOT_LIMIT_REACHED
  There is not enough space to create a snapshot.
- VOLUME_IS_NOT_CONSISTENT_SLAVE
  The operation not allowed on an inconsistent secondary volume.
- SNAPSHOT_GROUP_IS_INTERNAL
  Internal snapshots cannot be mapped, modified, or deleted.
- SNAPSHOT_GROUP_ILLEGAL_PRIORITY
  Illegal snapshot group priority; must be an integer between 1 and 4.
- SNAPSHOT_HAS_ACTIVE_SYNC_JOB
  The snapshot is currently the target of an active sync job.

  Troubleshooting: Please wait for the sync job to complete.
- OPERATION_DENIED_OBJECT_MANAGED
  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.
- CONS_GROUP_TOKEN_MISMATCH
  The token does not match the consistency group.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support
- MAX_SNAPSHOTS_PER_VOLUME_REACHED
  The maximum allowed number of snapshots is already reached.

Changing a snapshot group deletion priority

Use the **snap_group_change_priority** command to change the deletion priority of a snapshot group.

```
snap_group_change_priority snap_group=SnapshotGroupName  delete_priority=del_value
```

Parameters

Name	Type	Description	Mandatory
snap_group	Object name	Name of the snapshot group whose delete_priority is to be changed.	Y
delete_priority	Integer	Priority according to which this snapshot group is deleted.	Y

This command changes the priority of the deletion of an existing snapshot group. Similarly to snapshots, the system determines which of the snapshot groups is deleted first when it runs out of snapshot storage, in accordance with the redirect-on-write mechanism. When the system runs out of space, it deletes the snapshot or snapshot group with the highest deletion priority, and among them the unmapped snapshots or snapshot groups, and the snapshot or snapshot group which was created first.

See Changing a snapshot deletion priority for more details about the valid deletion priority values and their meaning.

Example:

```
snap_group_change_priority snap_group=DBbackup delete_priority=4
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission	Condition
Storage administrator	Allowed	N/A

User Category	Permission	Condition
Storage integration administrator	Allowed	N/A
Application administrator	Conditionally Allowed	At least one of the volumes in the master Consistency Group is mapped to a host or cluster associated with the user and Snapshot Group was created by a server administrator
Security administrator	Disallowed	N/A
Read-only users	Disallowed	N/A
Technicians	Disallowed	N/A

Return codes

- SNAPSHOT_GROUP_BAD_NAME
  The snapshot group name does not exist.
- SNAPSHOT_ILLEGAL_PRIORITY
  Illegal snapshot priority; must be an integer between 1 and 4.
- SNAPSHOT_GROUP_IS_INTERNAL
  Internal snapshots cannot be mapped, modified, or deleted.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support

Deleting a snapshot group

Use the **snap_group_delete** command to delete a snapshot group and all its snapshots.

```
snap_group_delete snap_group=SnapshotGroupName
```

Parameters

Name	Type	Description	Mandatory
snap_group	Object name	Name of the snapshot group to be deleted.	Y

This command deletes the snapshot group, as well as all of the snapshots that are contained in the snapshot group. Refer to the documentation on Deleting a snapshot for more information about deleting snapshots.

If one of the members of the snapshot group is mapped to a host, then the entire snapshot group cannot be deleted.

The command is inapplicable for a snapshot group that is still associated with a mirrored consistency group.

Example:

```
snap_group_delete snap_group=DBBackupweekly
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission	Condition
Storage administrator	Allowed	N/A
Storage integration administrator	Allowed	N/A
Application administrator	Conditionally Allowed	At least one of the volumes in the master Consistency Group is mapped to a host or cluster associated with the user and Snapshot Group was created by a server administrator
Security administrator	Disallowed	N/A
Read-only users	Disallowed	N/A
Technicians	Disallowed	N/A

Return codes

- OPERATION_DENIED_OBJECT_MANAGED
  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.
- SNAPSHOT_GROUP_BAD_NAME
  The snapshot group name does not exist.
- SNAPSHOT_IS_MAPPED
  A snapshot that is mapped to a host cannot be deleted.

- VOLUME_IS_BOUND
  The volume is bound to an ALU.

  Troubleshooting: Unbind the volume from the ALU.

- SNAPSHOT_HAS_ACTIVE_SYNC_JOB
  The snapshot is currently the target of an active sync job.

  Troubleshooting: Please wait for the sync job to complete.

- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support

Disbanding a snapshot group

Use the **snap_group_disband** command to disband a snapshot group into independent snapshots.

```
snap_group_disband snap_group=SnapshotGroupName
```

Parameters

Name	Type	Description	Mandatory
snap_group	Object name	Snapshot group to be disbanded.	Y

This command disbands the snapshot group into independent snapshots. After executing this command, the snapshots can be individually deleted, restored, unlocked, duplicated, and so on. The snapshot group does not exist anymore after this command. The snapshots retain the same names (snap_group_name.volumename).

The command is inapplicable for a snapshot group of a mirrored consistency group.

Example:

```
snap_group_disband snap_group=DBbackup_copy
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission	Condition
Storage administrator	Allowed	N/A
Storage integration administrator	Allowed	N/A
Application administrator	Conditionally Allowed	At least one of the volumes in the master Consistency Group is mapped to a host or cluster associated with the user and Snapshot Group was created by a server administrator
Security administrator	Disallowed	N/A
Read-only users	Disallowed	N/A
Technicians	Disallowed	N/A

Return codes

- OPERATION_DENIED_OBJECT_MANAGED
  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.

- SNAPSHOT_GROUP_BAD_NAME
  The snapshot group name does not exist.

- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support

Duplicating a snapshot group

Use the **snap_group_duplicate** command to duplicate an existing snapshot group.

```
snap_group_duplicate snap_group=SnapshotGroupName [ new_snap_group=NewName ]
```

Parameters

Name	Type	Description	Mandatory	Default
snap_group	Object name	Name of the snapshot group to be duplicated.	Y	N/A
new_snap_group	Object name	Name of the newly generated snapshot group.	N	Autogenerated name.

This command duplicates the specified snapshot group. This is functionally equivalent to duplicating all the snapshots in the snapshot group using Duplicating a snapshot and creating a new snapshot group that contains all the generated snapshots.

The name of the new snapshot group is either specified as a parameter or generated automatically.

Refer to Duplicating a snapshot for more details about the snapshot duplication operation.

Deletion priority:

- The deletion priority of the duplicated snapshots is 0.

Example:

```
snap_group_duplicate snap_group=DBbackup new_snap_group=DBbackup_copy
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission	Condition
Storage administrator	Allowed	N/A
Storage integration administrator	Allowed	N/A
Application administrator	Conditionally Allowed	At least one of the volumes in the master Consistency Group is mapped to a host or cluster associated with the user and Snapshot Group was created by a server administrator
Security administrator	Disallowed	N/A
Read-only users	Disallowed	N/A
Technicians	Disallowed	N/A

Return codes

- SNAPSHOT_GROUP_BAD_NAME
  The snapshot group name does not exist.
- MAX_VOLUMES_REACHED
  The maximum allowed number of volumes is already reached.
- DOMAIN_MAX_VOLUMES_REACHED
  The domain exceeds the maximum allowed number of volumes.
- SNAPSHOT_GROUP_NAME_EXISTS
  The snapshot group name already exists.
- OPERATION_DENIED_OBJECT_MANAGED
  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support
- MAX_SNAPSHOTS_PER_VOLUME_REACHED
  The maximum allowed number of snapshots is already reached.

Formatting a snapshot group

Use the **snap_group_format** command to format a snapshot group.

```
snap_group_format  snap_group=SnapshotGroupName
```

Parameters

Name	Type	Description	Mandatory
snap_group	Object name	The snapshot group to be formatted.	Y

This command deletes the content of a snapshot group while maintaining its snapshots mapping to the host. The format operation results with:

- The snapshots of the formatted snapshot group are read-only
- The format operation has no impact on performance
- The snapshots of the formatted snapshot group do not consume space
- Reading from the snapshots of the formatted snapshot group always returns zeroes

- The snapshots can be overridden
- The snapshots can be deleted
- The snapshots deletion priority can be changed

Example:

```
snap_group_format snap_group
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission	Condition
Storage administrator	Allowed	N/A
Storage integration administrator	Allowed	N/A
Application administrator	Conditionally Allowed	At least one of the volumes in the master Consistency Group is mapped to a host or cluster associated with the user and Snapshot Group was created by a server administrator
Security administrator	Disallowed	N/A
Read-only users	Disallowed	N/A
Technicians	Disallowed	N/A

Return codes

- SNAPSHOT_HAS_ACTIVE_SYNC_JOB
  The snapshot is currently the target of an active sync job.

  Troubleshooting: Please wait for the sync job to complete.
- SNAPSHOT_GROUP_IS_FORMATTED
  The snapshot group is formatted.
- ELCS_GROUP_CANNOT_BE_FORMATTED
  The snapshot group is an ELCS (external last consistent snapshot), and cannot be formatted.
- MAX_VOLUMES_REACHED
  The maximum allowed number of volumes is already reached.
- SNAPSHOT_GROUP_IS_INTERNAL
  Internal snapshots cannot be mapped, modified, or deleted.
- VOLUME_IS_NOT_A_SNAPSHOT
  The operation is permitted on snapshots only.
- SNAPSHOT_GROUP_BAD_NAME
  The snapshot group name does not exist.
- OPERATION_DENIED_OBJECT_MANAGED
  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support

Listing snapshot groups

Use the **snap_group_list** command to list all snapshot groups or a specific one.

```
snap_group_list [ snap_group=SnapshotGroupName | cg=cgName ] [ managed=<yes|no|all> ]
```

Parameters

Name	Type	Description	Mandatory	Default
snap_group	Object name	Name of a specific snapshot group to be listed.	N	All snapshot groups.
cg	Object name	List all the snapshot groups of this Consistency Group.	N	All snapshot groups.
managed	Boolean	Defines whether to show unmanaged snap groups (no), managed (yes) or both (all).	N	no.

This command lists snapshot groups. When a snapshot group name is specified, then only that specific snapshot group is listed. When a consistency group name is specified, then the snapshot groups of this consistency group are listed.

This command displays the following snapshot group format field (available on the XML output format):

- snap_group_format

Field ID	Field output	Default position
name	Name	1
cg	CG	2
snapshot_time	Snapshot Time	3
locked	Locked	N/A
modified	Modified	N/A
delete_priority	Deletion Priority	4
snap_group_format	Snapshot Group Format	N/A
snap_group_descriptor	Snapshot Group Descriptor	N/A
managed	Managed	N/A

Example:

```
snap_group_list cg=DBvolumes
```

Output:

```
Name            CG          Snapshot Time       Deletion Priority
DBbackup        DBvolumes   2007-01-03 17:46:29 1
DBbackupdaily   DBvolumes   2007-01-03 17:49:36 1
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Locking a snapshot group

Use the **snap_group_lock** command to lock a snapshot group by locking all its snapshots.

```
snap_group_lock snap_group=SnapshotGroupName
```

Parameters

Name	Type	Description	Mandatory
snap_group	Object name	Name of the snapshot group to be locked.	Y

This command is functionally equivalent to locking all snapshots individually (through executing Locking a volume on each snapshot). Refer to the documentation of Locking a volume for a description of locking behavior.

Example:

```
snap_group_lock snap_group=DBbackup
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission	Condition
Storage administrator	Allowed	N/A
Storage integration administrator	Allowed	N/A
Application administrator	Conditionally Allowed	At least one of the volumes in the master Consistency Group is mapped to a host or cluster associated with the user and Snapshot Group was created by a server administrator
Security administrator	Disallowed	N/A
Read-only users	Disallowed	N/A
Technicians	Disallowed	N/A

Return codes

- SNAPSHOT_GROUP_BAD_NAME
  The snapshot group name does not exist.
- SNAPSHOT_GROUP_IS_INTERNAL

Internal snapshots cannot be mapped, modified, or deleted.

- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support

Renaming a snapshot group

Use the **snap_group_rename** command to rename a snapshot group.

```
snap_group_rename snap_group=SnapshotGroupName new_name=Name
```

Parameters

Name	Type	Description	Mandatory
snap_group	Object name	Name of the snapshot group to be renamed.	Y
new_name	Object name	New name for the snapshot group.	Y

Example:

```
snap_group_rename snap_group=DBbackup new_name=DBBackupweekly
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission	Condition
Storage administrator	Allowed	N/A
Storage integration administrator	Allowed	N/A
Application administrator	Conditionally Allowed	At least one of the volumes in the master Consistency Group is mapped to a host or cluster associated with the user and Snapshot Group was created by a server administrator
Security administrator	Disallowed	N/A
Read-only users	Disallowed	N/A
Technicians	Disallowed	N/A

Return codes

- SNAPSHOT_GROUP_BAD_NAME
  The snapshot group name does not exist.
- SNAPSHOT_GROUP_NAME_EXISTS
  The snapshot group name already exists.
- OPERATION_DENIED_OBJECT_MANAGED
  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support

Restoring a consistency group from a snapshot group

Use the **snap_group_restore** command to restore the master volumes of a consistency group, or of a snapshot group from one of its associated snapshot groups.

```
snap_group_restore snap_group=SnapshotGroupName [ target_snap_group=SnapshotGroupName ]
```

Parameters

Name	Type	Description	Mandatory
snap_group	Object name	Name of the snapshot group from which to restore its master volumes.	Y
target_snap_group	Object name	Snapshot group to be restored.	N

Using this command is equivalent to restoring all the volumes in the consistency group, or all the snapshots in the target snapshot group from their snapshots in the snapshot group.

It is possible to restore a snapshot group from a snapshot group.

Requirements for a successful command completion:

- The consistency group or the target snapshot group must contain the exact same volumes that they contained when the snapshot group was generated.
  - Each volume added to the consistency group after the creation of the snapshot group must be removed from the consistency group before restoration is completed.
- The command is inapplicable for a snapshot group of a mirrored consistency group.

See Restoring a volume from a snapshot for more information about the restoring.

Example:

```
snap_group_restore snap_group=DBbackup_copy
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission	Condition
Storage administrator	Allowed	N/A
Storage integration administrator	Allowed	N/A
Application administrator	Conditionally Allowed	Both target and source are snapshots groups of the same master Consistency Group, where at least one of the master volumes in this Consistency Group is mapped to a host or cluster associated with the user, and the target Snapshot Group was created by an application administrator.
Security administrator	Disallowed	N/A
Read-only users	Disallowed	N/A
Technicians	Disallowed	N/A

Return codes

- VOLUME_HAS_DATA_MIGRATION
  Data Migration is defined for this volume.
- SNAPSHOT_GROUP_BAD_NAME
  The snapshot group name does not exist.
- CONS_GROUP_MISMATCH
  The snapshot group does not match the consistency group volumes.
- VOLUME_TOO_BIG
  No space to allocate to the volume.
- VOLUME_HAS_MIRROR
  A mirror is defined for this volume.
- VOLUME_HAS_HA
  This operation is forbidden on a volume with a HyperSwap relation.
- CONS_GROUP_HAS_MIRROR
  Mirroring is defined for this consistency group.
- VOLUME_LOCKED
  The volume is locked.
- TARGET_SNAPSHOT_GROUP_BAD_NAME
  The target snapshot group name does not exist.
- SNAPSHOT_GROUP_MISMATCH
  The snapshot group does not match the target snapshot group.
- TARGET_SNAPSHOT_GROUP_SAME_AS_SOURCE
  The target snapshot group is identical with the snapshot group.
- OPERATION_DENIED_OBJECT_MANAGED
  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support

Unlocking a snapshot group

Use the **snap_group_unlock** command to unlock a snapshot group by unlocking all its snapshots.

```
snap_group_unlock snap_group=SnapshotGroupName
```

Parameters

Name	Type	Description	Mandatory
snap_group	Object name	Name of the snapshot group to be unlocked.	Y

This command unlocks a snapshot group by unlocking all its snapshots. This is equivalent to executing Unlocking a volume on each snapshot. Refer to the documentation of Unlocking a volume for a description of unlocking behavior.

Example:

```
snap_group_unlock snap_group=DBbackup
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission	Condition
Storage administrator	Allowed	N/A
Storage integration administrator	Allowed	N/A
Application administrator	Conditionally Allowed	At least one of the volumes in the master Consistency Group is mapped to a host or cluster associated with the user and Snapshot Group was created by a server administrator
Security administrator	Disallowed	N/A
Read-only users	Disallowed	N/A
Technicians	Disallowed	N/A

Return codes

- SNAPSHOT_GROUP_BAD_NAME
  The snapshot group name does not exist.
- SNAPSHOT_GROUP_IS_INTERNAL
  Internal snapshots cannot be mapped, modified, or deleted.
- OPERATION_DENIED_OBJECT_MANAGED
  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support

Setting a snapshot group descriptor

Use the **snap_group_set_descriptor** command to set a snapshot group descriptor.

```
snap_group_set_descriptor snap_group=SnapshotGroupName descriptor=Descriptor
```

Parameters

Name	Type	Description	Mandatory
snap_group	Object name	Name of the snapshot group.	Y
descriptor	String	A snap group descriptor to be used by external software.	Y

Provides external software with the ability to mark the snapshot as part of a consistency group for various usage scenarios. The command replaces an existing descriptor with a newly specified one.

Example:

```
snap_group_set_descriptor snap_group=DBbackup descriptor=blabla
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission	Condition
Storage administrator	Allowed	N/A

User Category	Permission	Condition
Storage integration administrator	Allowed	N/A
Application administrator	Conditionally Allowed	At least one of the volumes in the master Consistency Group is mapped to a host or cluster associated with the user and Snapshot Group was created by a server administrator
Security administrator	Disallowed	N/A
Read-only users	Disallowed	N/A
Technicians	Disallowed	N/A

Return codes

- SNAPSHOT_GROUP_BAD_NAME
  The snapshot group name does not exist.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support

Returning a snapshot group's descriptor

Use the **snap_group_get_descriptor** command to return a snapshot group's descriptor.

```
snap_group_get_descriptor snap_group=SnapshotGroupName
```

Parameters

Name	Type	Description	Mandatory
snap_group	Object name	Name of the snapshot group.	Y

The command provides an external software with the ability to obtain the descriptor attribute value for a snapshot group.

Example:

```
snap_group_get_descriptor snap_group=DBbackup
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Return codes

- SNAPSHOT_GROUP_BAD_NAME
  The snapshot group name does not exist.

Storage pool management commands

This section describes the command-line interface (CLI) for storage pool management.

See also:

- Volume management commands
- Volume snapshot management commands
- Consistency group management commands

- **Moving a consistency group between storage pools**
  Use the **cg_move** command to move a consistency group, all its volumes, and all their snapshots and snapshot sets from one storage pool to another.
- **Changing the pool limitation, performance class, or threshold parameters**
  Use the **pool_change_config** command to change a storage pool configuration.

- **Changing pool settings for snapshots**
  Use the **pool_config_snapshots** command to change storage pool snapshot settings.
- **Creating storage pools**
  Use the **pool_create** command to create a storage pool.
- **Deleting a storage pool**
  Use the **pool_delete** command to delete a storage pool.
- **Listing storage pools**
  Use the **pool_list** command to list all storage pools or the specified one.
- **Renaming a storage pool**
  Use the **pool_rename** command to rename the specified storage pool.
- **Resizing a storage pool**
  Use the **pool_resize** command to resize a storage pool.
- **Moving a volume between storage pools**
  Use the **vol_move** command to move a volume and all its snapshot from one storage pool to another.

Moving a consistency group between storage pools

Use the **cg_move** command to move a consistency group, all its volumes, and all their snapshots and snapshot sets from one storage pool to another.

```
cg_move cg=cgName pool=PoolName [ domain_adjust=<yes|no> ]
```

Parameters

Name	Type	Description	Mandatory	Default
cg	Object name	Name of the consistency group to be moved.	Y	N/A
pool	Object name	Name of the target storage pool.	Y	N/A
domain_adjust	Boolean	Adjusts domain resources. If set to True, the resources of the consistency group source domain and destination domain are adjusted to accommodate the consistency group being moved.	N	no

For successful command completion, there must be sufficient space on the target pools. If the consistency group is mirrored, it can only be moved to a pool that is not thin-provisioned.

Example:

```
cg_move cg=DBGroup pool=DBPool
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- CONS_GROUP_BAD_NAME
  The consistency group name does not exist.
- POOL_DOES_NOT_EXIST
  The storage pool does not exist.
- NOT_ENOUGH_SPACE
  No space to allocate for the volume's current usage.
- NOT_ENOUGH_SNAPSHOT_SPACE
  Snapshot usage will exceed the snapshot limit.
- DOMAIN_MAX_VOLUMES_REACHED
  The domain exceeds the maximum allowed number of volumes.
- MAX_VOLUMES_REACHED
  The maximum allowed number of volumes is already reached.
- DOMAIN_MAX_CONS_GROUPS_REACHED
  The domain exceeds the maximum allowed number of consistency groups.

- MAX_CONS_GROUPS_REACHED
  The maximum allowed number of consistency groups is already reached.
- DOMAIN_MAX_MIRRORS_REACHED
  The domain exceeds the maximum allowed number of mirrors.
- MAX_MIRRORS_REACHED
  The maximum number of mirrors is already reached.
- DOMAIN_USED_TARGET_NOT_IN_DESTINATION
  A target that is used by mirror in the pool is not associated with the target domain.
- DOMAIN_USED_SCHEDULE_NOT_IN_DESTINATION
  A schedule used by a mirror in the pool is not associated with the target domain.
- MAPPED_HOSTS_NOT_IN_DESTINATION
  A host that is mapped to a volume in the pool is not associated with the target domain.
- MAPPED_CLUSTERS_NOT_IN_DESTINATION
  A cluster that is mapped to a volume in the pool is not associated with the target domain.
- OPERATION_DENIED_OBJECT_MANAGED
  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.
- CONS_GROUP_REQUIRES_DESTINATION_POOL
  A destination pool must be defined.
- MAX_DMS_REACHED
  The maximum number of remote volumes (mirror/migration) is already reached.

  Troubleshooting: Delete unnecessary data migration objects.
- DOMAIN_MAX_DMS_REACHED
  The domain exceeds the maximum allowed number of data migrations.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support
- NO_SPACE
  The system does not have enough free space for the requested storage pool size.
- VOLUME_TOO_BIG
  No space to allocate to the volume.

Changing the pool limitation, performance class, or threshold parameters

Use the **pool_change_config** command to change a storage pool configuration.

```
pool_change_config pool=PoolName [ lock_behavior=<read_only|no_io> ] [ perf_class=perfClassName ]
 [ restore_thresholds=<yes|no> | hysteresis=HysteresisValue | < code=EventCode severity=
<INFORMATIONAL|WARNING|MINOR|MAJOR|CRITICAL|NONE> threshold=<ThresholdValue|NONE> > ]
```

Parameters

Name	Type	Description	Mandatory	Default
pool	Object name	The name of a storage pool.	Y	N/A
lock_behavior	Enumeration	Determines whether and how the pool is locked upon space depletion.	N	read_only
perf_class	Object name	The name of the performance class pool.	N	No performance class
code	N/A	Event code.	N	No code
severity	Enumeration	Severity.	N	No severity
threshold	Integer	The threshold value. None indicates that an event with this severity is not created.	N	No threshold
restore_thresholds	Boolean	Restore thresholds to default values.	N	no
hysteresis	Integer	The hysteresis of the event throwing.	N	"3"

This command changes the pool behavior when the pool runs out of thin provisioning space.

For thin provisioned storage pools, the lock_behavior parameter sets how the pool is locked upon space depletion. The pool can be locked for write, or for both read and write.

Example:

```
pool_change_config pool=VOL_BREED_None_0 lock_behavior=read_only
```

This command changes the Performance Class of the pool.

Example:

```
pool_change_config pool=VOL_BREED_None_1 perf_class=valid_perf_class_name
```

This command changes the thresholds parameters of the pool or reset it to default thresholds value.

Example:

```
pool_change_config pool=VOL_BREED_None_1 code=STORAGE_POOL_VOLUME_USAGE_INCREASED severity=INFORMATIONAL threshold=40
pool_change_config pool=VOL_BREED_None_1 code=STORAGE_POOL_SNAPSHOT_USAGE_INCREASED severity=INFORMATIONAL threshold=50
pool_change_config pool=VOL_BREED_None_1 restore_thresholds=yes
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- POOL_DOES_NOT_EXIST
  The storage pool does not exist.
- PERF_CLASS_BAD_NAME
  The performance class does not exist.
- PERF_CLASS_ASSOCIATED_WITH_HOSTS
  Performance class *Performance Class* is already being used by a host.
- PERF_CLASS_ASSOCIATED_WITH_VOLUMES
  Performance class *Performance Class* is already being used by a volume.
- POOL_ALREADY_IN_PERF_CLASS
  Pool *pool name* is already in performance class *Performance Class*.
- OPERATION_DENIED_OBJECT_MANAGED
  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.
- UNRECOGNIZED_EVENT_CODE
  '*String*' is not a recognized return code.

  Troubleshooting: Consult the manual for the list of valid return codes.
- EVENT_DOES_NOT_HAVE_THRESHOLDS
  The event does not have thresholds.
- EVENT_THRESHOLD_IS_ILLEGAL
  An illegal value for the event threshold.

  Troubleshooting: Event threshold values must be monotonic.

Changing pool settings for snapshots

Use the **pool_config_snapshots** command to change storage pool snapshot settings.

```
pool_config_snapshots pool=PoolName [ protected_snapshot_priority=<0|1|2|3|4> ]
```

Parameters

Name	Type	Description	Mandatory	Default
pool	Object name	The name of a storage pool.	Y	N/A
protected_snapshot_priority	Integer	Specifies the snapshot delete priority from 0 to 4 (see full explanation below).	N	unchanged

This command changes the storage pool snapshot limitation policy.

The *create_last_consistent_snapshot* attribute (used for systems which have no space):

- If the value of the attribute is No, no last consistent snapshot is generated.
- If the value is changed while synchronizing, the existing snapshot is not deleted.

The protected_snapshot_priority parameter:

- Snapshots with a lower delete priority (that is, a higher number) than the specified value might be deleted by the system automatically, in order to free space, before pausing the mirroring, thus protecting snapshots with a priority equal or higher than the value.
- If, for example, the value is set to 3:
  - The system will deactivate mirroring if not enough space can be freed even after the deletion of snapshots with deletion priority of 4.
  - Snapshots with priority level 1, 2 and 3 will not be deleted.
- If the value is set to 4, the system will deactivate mirroring before deleting any of the snapshots.
- If the value is set to 0, the system can delete any snapshot regardless of deletion priority.

Example:

```
pool_config_snapshots pool=DBPool
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_CHANGE_THE_PROTECTED_LEVEL_OF_SNAPSHOTS
  Are you sure you want to change the protection level of a snapshot in storage pool *Pool*? Note that in case of pool space depletion the system will delete protected snapshots only after deleting unprotected snapshots and internal asynchronous mirror snapshots.
- ARE_YOU_SURE_YOU_WANT_TO_INCREASE_THE_PROTECTED_LEVEL_OF_EXISTING_SNAPSHOTS
  Are you sure you want to increase the protection level of a snapshot in storage pool *Pool*? Note that the pool contains unprotected snapshots that will become protected after issuing this command. In case of pool space depletion the system will delete protected snapshots only after deleting unprotected snapshots and internal asynchronous mirror snapshots.
- ARE_YOU_SURE_YOU_WANT_TO_DECREASE_THE_PROTECTED_LEVEL_OF_EXISTING_SNAPSHOTS
  Are you sure you want to decrease the protection level of a snapshot in Storage Pool *Pool*? Note that the pool contains protected snapshots that will become unprotected after issuing this command. In case of pool space depletion the system will delete internal asynchronous mirror snapshots only after deleting unprotected snapshots.

Return codes

- POOL_DOES_NOT_EXIST
  The storage pool does not exist.
- SNAPSHOT_ILLEGAL_PRIORITY
  Illegal snapshot priority; must be an integer between 1 and 4.
- OPERATION_DENIED_OBJECT_MANAGED
  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.

Creating storage pools

Use the **pool_create** command to create a storage pool.

```
pool_create pool=PoolName size=GB snapshot_size=GB [ lock_behavior=<read_only|no_io> ]
 [ perf_class=perfClassName ] [ domain=DomainName ]
```

Parameters

Name	Type	Description	Mandatory	Default
pool	Object name	The name of the new storage pool.	Y	N/A
size	Positive integer	Effective capacity of the storage pool (in gigabytes).	Y	N/A
snapshot_size	Positive integer	Effective capacity allocated for snapshots.	Y	N/A
lock_behavior	Enumeration	Determines whether and how the pool is locked upon space depletion.	N	read_only

Name	Type	Description	Mandatory	Default
perf_class	Object name	The name of the performance class pool.	N	No performance class
domain	Object name	Add the pool to the specified domain.	N	none

The name of the storage pool must be unique in the system. Upon creation, the storage pool is empty and does not contain volumes.

Pool size limits

The parameters size and snapshot_size relate to effective capacity.

The upper limit of the size parameter is not restricted, provided that the total effective capacity of all pools does not exceed the storage system's maximum effective capacity. However, to protect the system against overallocation of its resources, the maximum allowed total effective capacity of the volumes in a consistency group is limited to the effective capacity of the system.

As for the lower limits of the size and snapshot_size parameters, the following restrictions apply:

- size cannot be less than 2 TB
- snapshot_size must be 0, or not less than 400 GB.

These limits can be ignored by force (-y). To change the limits, contact the IBM Support.

Example:

```
pool_create pool=DBPool size=1000 snapshot_size=500
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- POOL_SIZE_SMALL
  The pool size is very small. Volumes may not be able to use this space efficiently. Are you sure?
- POOL_SNAPSHOT_SIZE_SMALL
  The pool snapshot size is very small. Snapshots may be deleted frequently. Are you sure?

Return codes

- POOL_NAME_EXISTS
  The storage pool name already is assigned to another storage pool.
- PERF_CLASS_BAD_NAME
  The performance class does not exist.
- PERF_CLASS_ASSOCIATED_WITH_HOSTS
  Performance class *Performance Class* is already being used by a host.
- PERF_CLASS_ASSOCIATED_WITH_VOLUMES
  Performance class *Performance Class* is already being used by a volume.
- MAX_POOLS_REACHED
  The maximum allowed number of storage pools is already reached.
- NO_SPACE
  The system does not have enough free space for the requested storage pool size.
- SNAPSHOT_SIZE_BIGGER_THAN_POOL_SIZE
  The snapshot size must be equal to or smaller than the pool size.
- REACHED_POOL_MAX_SIZE
  Maximum pool size usage is already reached.
- DOMAIN_DOESNT_EXIST
  The domain does not exist.
- USER_ASSOCIATED_TO_MORE_THAN_ONE_DOMAIN
  The current user is attached to more than one domain, and it is not clear in which domain the pool is to be created.

Troubleshooting: Re-run the command by specifying a domain.

- NO_FREE_CAPACITY_IN_DOMAIN
  There is not enough free space in the domain.
- DOMAIN_MAX_POOLS_REACHED
  The maximum allowed number of domain pools is already reached.

Deleting a storage pool

Use the **pool_delete** command to delete a storage pool.

```
pool_delete pool=PoolName
```

Parameters

Name	Type	Description	Mandatory
pool	Object name	The name of the storage pool to be deleted.	Y

This command fails if the storage pool is not empty, that is it still contains volumes.

The capacity of the deleted storage pool is added to the free space.

Example:

```
pool_delete pool=ERPPool
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_DELETE_POOL
  Are you sure you want to delete storage pool *Pool*?

Return codes

- POOL_DOES_NOT_EXIST
  The storage pool does not exist.
- POOL_HAS_CG
  The storage pool comprises consistency groups.
- POOL_IN_USE
  The storage pool comprises allocated volumes.
- OPERATION_DENIED_OBJECT_MANAGED
  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.

Listing storage pools

Use the **pool_list** command to list all storage pools or the specified one.

```
pool_list [ pool=PoolName ] [ managed=<yes|no|all> ] [ domain=DomainName ]
```

Parameters

Name	Type	Description	Mandatory	Default
pool	Object name	The name of a storage pool.	N	All pools.
managed	Boolean	Determines whether to show unmanaged pools (no), managed (yes), or both (all).	N	No

Name	Type	Description	Mandatory	Default
domain	Object name	The domain name.	N	All Domains

When the pool parameter is provided, only the specified storage pool is listed.

Example:

```
pool_list
```

Output:

```
Name     Size (GB)  Empty Space (GB)
default  24292      9225
DBPool   1013       1013
```

Field ID	Field output	Default position
name	Name	1
size	Size (GB)	2
size_MiB	Size (MiB)	N/A
snapshot_size	Snap Size (GB)	3
snapshot_size_MiB	Snap Size (MiB)	N/A
total_volume_size	Total Vols (GB)	4
total_volume_size_MiB	Total Vols (MiB)	N/A
empty_space	Empty (GB)	5
empty_space_MiB	Empty (MiB)	N/A
used_by_volumes	Used by Vols (GB)	6
used_by_volumes_MiB	Used by Vols (MiB)	N/A
used_by_snapshots	Used by Snaps (GB)	7
used_by_snapshots_MiB	Used by Snaps (MiB)	N/A
creator	Creator	N/A
locked	Locked	8
lock_behavior	Lock Behavior	N/A
create_last_consistent_snapshot	Create Last Consistent Snapshot	N/A
protected_snapshot_priority	Protected Snapshots Priority	N/A
managed	Managed	N/A
perf_class	Perf Class Name	9
domain	Domain	10
sparse	Sparse	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Renaming a storage pool

Use the **pool_rename** command to rename the specified storage pool.

```
pool_rename pool=PoolName new_name=Name
```

Parameters

Name	Type	Description	Mandatory
pool	Object name	The current name of the storage pool.	Y
new_name	Object name	The new name of the storage pool.	Y

The new name of the storage pool must be unique in the system.

This command succeeds even if the new name is identical with the current name.

Example:

```
pool_rename pool=DBPool new_name=ERPPool
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- POOL_DOES_NOT_EXIST
  The storage pool does not exist.
- POOL_NAME_EXISTS
  The storage pool name already is assigned to another storage pool.
- OPERATION_DENIED_OBJECT_MANAGED
  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.

Resizing a storage pool

Use the **pool_resize** command to resize a storage pool.

```
pool_resize pool=PoolName [ size=GB ] [ snapshot_size=GB ]
```

Parameters

Name	Type	Description	Mandatory	Default
pool	Object name	The name of the storage pool to be resized.	Y	N/A
size	Positive integer	The new size of the storage pool (in gigabytes)	N	N/A
snapshot_size	Integer	The new limit on snapshot capacity usage of the storage pool.	N	Leave unchanged.

The command can either increase or decrease the storage pool size.

The new size of the storage pool is specified as an integer multiple of 10^9 bytes, but the actual size of the created storage pool is rounded up to the nearest integer multiple of $16x2^{30}$ bytes.

Capacity accounting is performed relative to the free space.

You can either specify both hard size and soft size or one size only, which assumes that the hard size and the soft size are identical.

- When increasing a storage pool size, the command succeeds only if the free space holds enough free capacity to allow such an increase.
- When decreasing a storage pool size, the command succeeds only if the storage pool itself holds enough free capacity to allow such a reduction.
- If the new size equals the current size, the command succeeds without changing the storage pool.

Pool size limits

The parameters size and snapshot_size relate to effective capacity.

The upper of the size parameter is set to the effective system capacity.

As for the lower limits of the size and snapshot_size parameters, the following restrictions apply:

- size cannot be less than 2 TB
- snapshot_size must be 0, or not less than 400 GB.

These limits can be ignored by force (-y). To change the limits, contact the IBM Support.

This command fails if the current storage pool size (hard or soft) cannot be decreased, or if the free space (hard or soft) cannot be decreased.

Resizing a pool that stores asynchronously mirrored volumes or consistency groups:

- The pool's hard and soft sizes must be identical.

Example:

```
pool_resize pool=DBPool size=1300
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- POOL_SIZE_SMALL
  The pool size is very small. Volumes may not be able to use this space efficiently. Are you sure?
- POOL_SNAPSHOT_SIZE_SMALL
  The pool snapshot size is very small. Snapshots may be deleted frequently. Are you sure?

Return codes

- POOL_DOES_NOT_EXIST
  The storage pool does not exist.
- POOL_SIZE_TOO_SMALL
  Storage pool usage exceeds the requested size.
- REACHED_POOL_MAX_SIZE
  Maximum pool size usage is already reached.
- NO_SPACE
  The system does not have enough free space for the requested storage pool size.
- POOL_SNAPSHOT_SIZE_TOO_SMALL
  Storage pool snapshot usage exceeds the requested snapshot size.
- SNAPSHOT_SIZE_BIGGER_THAN_POOL_SIZE
  The snapshot size must be equal to or smaller than the pool size.
- OPERATION_DENIED_OBJECT_MANAGED
  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.
- NO_FREE_CAPACITY_IN_DOMAIN
  There is not enough free space in the domain.

Moving a volume between storage pools

Use the **vol_move** command to move a volume and all its snapshot from one storage pool to another.

```
vol_move vol=VolName pool=PoolName [ domain_adjust=<yes|no> ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	Name of the volume to move.	Y	N/A
pool	Object name	Name of the storage pool to which to move.	Y	N/A
domain_adjust	Boolean	Adjust domain resources. If set to *true*, the resources of the volume source domain and destination domain are adjusted to accommodate the volume being moved.	N	no

When moving a master volume from one storage pool to another, all of its snapshots are moved together with it to the destination storage pool.

This command fails when trying to move a snapshot of a volume on its own. This command can fail due to the lack of either soft or hard space.

The command succeeds only if the destination storage pool has enough free storage capacity to accommodate the volume and its snapshots. The exact amount of storage capacity allocated from the destination storage pool is released at the source storage pool.

A volume which belongs to a consistency group cannot be moved without the entire consistency group. You may use Moving a consistency group between storage pools to move the consistency group itself from one storage pool to another.

A volume that is asynchronously mirrored cannot be moved into a thin provisioning pool.

Example:

```
vol_move vol=DBLog pool=DBPool
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.
- POOL_DOES_NOT_EXIST
  The storage pool does not exist.
- NOT_ENOUGH_SPACE
  No space to allocate for the volume's current usage.
- NO_SPACE
  The system does not have enough free space for the requested storage pool size.
- VOLUME_TOO_BIG
  No space to allocate to the volume.
- NO_FREE_CAPACITY_IN_DOMAIN
  There is not enough free space in the domain.
- VOLUME_IS_SNAPSHOT
  The operation is not permitted on snapshots.
- VOLUME_HAS_OLVM
  An IBM Hyper-Scale Mobility relationship is defined for this volume.
- VOLUME_BELONGS_TO_CG
  The volume belongs to a consistency group.
- NOT_ENOUGH_SNAPSHOT_SPACE
  Snapshot usage will exceed the snapshot limit.
- OPERATION_DENIED_OBJECT_MANAGED
  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.
- MAPPED_HOSTS_NOT_IN_DESTINATION
  A host that is mapped to a volume in the pool is not associated with the target domain.
- MAPPED_CLUSTERS_NOT_IN_DESTINATION
  A cluster that is mapped to a volume in the pool is not associated with the target domain.
- DOMAIN_USED_SCHEDULE_NOT_IN_DESTINATION
  A schedule used by a mirror in the pool is not associated with the target domain.
- DOMAIN_USED_TARGET_NOT_IN_DESTINATION
  A target that is used by mirror in the pool is not associated with the target domain.
- DOMAIN_MAX_MIRRORS_REACHED
  The domain exceeds the maximum allowed number of mirrors.
- DOMAIN_MAX_DMS_REACHED
  The domain exceeds the maximum allowed number of data migrations.
- DOMAIN_MAX_VOLUMES_REACHED
  The domain exceeds the maximum allowed number of volumes.
- MAX_MIRRORS_REACHED
  The maximum number of mirrors is already reached.
- MAX_DMS_REACHED
  The maximum number of remote volumes (mirror/migration) is already reached.

  Troubleshooting: Delete unnecessary data migration objects.
- MAX_VOLUMES_REACHED
  The maximum allowed number of volumes is already reached.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

Troubleshooting: Contact IBM Support

System management commands

This section describes the command-line interface (CLI) for system management.

- **Displaying the values of configuration parameters**
  Use the **config_get** command to show the values of configuration parameters.
- **Setting configuration parameters**
  Use the **config_set** command to set configuration parameters.
- **Testing the DNS**
  Use the **dns_test** command to test the DNS (Domain Naming Service).
- **Displaying help**
  Use the **help** command to display system help.
- **Printing the current maintenance urgency**
  Use the **maintenance_urgency_list** command to display the current maintenance urgency of the system.
- **Shutting down the system**
  Use the **shutdown** command to shut down the system.
- **Listing the operational state**
  Use the **state_list** command to display the current operational state of the system.
- **Displaying system usage and data reduction statistics**
  Use the **system_usage_get** command to retrieve system-wide usage and data reduction related statistics.
- **Displaying information about physical and effective capacity**
  Use the **system_capacity_list** command to display information about the system's effective and physical capacity.
- **Displaying information about effective capacity**
  Use the **system_effective_capacity_get** command to display additional information about the system's effective capacity.
- **Displaying system capacity thresholds**
  Use the **system_capacity_threshold_list** command to list the current system capacity thresholds.
- **Changing a system capacity threshold**
  Use the **system_capacity_threshold_change** command to change a system capacity threshold.
- **Resuming the system's normal operation**
  Use the **system_resume_normal_operation** command to move a system back to read-write state after it ran out of physical space.
- **Displaying the current time**
  Use the **time_list** command to display the current system time.
- **Setting the system's time**
  Use the **time_set** command to set the system's time in YYYY-MM-DD.HH:MM:SS format.
- **Listing optional time zones**
  Use the **timezone_list** command to list all optional time zones.
- **Setting the time zone**
  Use the **timezone_set** command to set the time zone of the system.
- **Printing the current system version**
  Use the **version_get** command to print the current version of the system.
- **Displaying the values of VPD parameters**
  Use the **vpd_config_get** command to display the values of VPD parameters.
- **Setting VPD parameters**
  Use the **vpd_config_set** command to set the values of VPD (Vital Product Data) parameters.
- **Displaying the system's MIB file**
  Use the **mib_get** command to display the system's MIB file.
- **Retrieving the electronic license acceptance status**
  Use the **elicense_status_get** command to retrieve the electronic license acceptance status.
- **Retrieving a fragment of the electronic license file**
  Use the **elicense_blob_get** command to retrieve a fragment of the electronic license file.
- **Accepting the electronic license agreement**
  Use the **elicense_accept** command to accept the electronic license agreement.
- **Enabling command auditing**
  Use the **audit_enable** command to enable CLI command auditing
- **Disabling command auditing**
  Use the **audit_disable** command to disable CLI command auditing.
- **Displaying the command audit state**
  Use the **audit_show** command to show the current state of CLI command auditing.
- **Configuring audit servers**
  Use the **audit_config_set** command to configure CLI command auditing.
- **Checking the command audit state**
  Use the **audit_config_get** command to show the current configuration of CLI command auditing.
- **Retrieving the list of Flash VDisks**
  Use the **flash_vdisk_list** command to retrieve the list of Flash VDisks
- **Enabling CIM service**
  Use the **cim_enable** command to enable the CIM service.
- **Disabling the CIM service**
  Use the **cim_disable** command to disable the CIM service.
- **Displaying the CIM service state**
  Use the **cim_show** command to display the current state of CIM service.

Displaying the values of configuration parameters

Use the **config_get** command to show the values of configuration parameters.

```
config_get [ name=Name ]
```

Parameters

Name	Type	Description	Mandatory	Default
name	String	Name of parameter to print.	N	All parameters.

Field ID	Field output	Default position
name	Name	1
value	Value	2

This command shows the name and value of the specified configuration parameter or of all of them, if no parameter is provided.

The values of the following parameters can be shown:

- dns_primary - IP address of the master DNS server.
- dns_secondary - IP address of the slave DNS server.
- email_reply_to_address - Reply-to address to be used when sending emails. This is useful for troubleshooting errors in email addresses.
- email_sender_address - Email address used as the sender's address when sending email messages.
- email_subject_format - Controls the formatting of the email subject line. To insert the event's data, use the following tags: {severity}, {description}, or {system_name}. System default is "{severity}: {description}".
- iscsi_name - iSCSI initiator name. Used when configuring a non-XIV system for data migration over iSCSI.
- machine_model
- machine_serial_number
- machine_type
- ntp_server - IP address or DNS name of the NTP server.
- snmp_community - Community used for SNMP queries of the system.
- snmp_location - SNMP location as shown in the SNMP MIB. (.1.3.6.1.2.1.1.6.0).
- snmp_contact - SNMP contact as shown in the SNMP MIB. (.1.3.6.1.2.1.1.4.0).
- snmp_trap_community - Community used for SNMP traps sent by the system.
- snmp_trap_type - The valid options are: common, severity_specific.
- snmp_type - SNMP version the snmpd should work with. The valid options are: none, v2c, v3, both.
- support_center_port_type - The valid options are: Management, VPN.
- system_id - Unique system identifier (equivalent to a serial number).
- system_name
- snmpv3_user
- snmpv3_encryption_type - The valid options are: DES, AES.
- snmpv3_encryption_passphrase
- snmpv3_authentication_type - The valid options are: MD5, SHA.
- snmpv3_authentication_passphrase
- ipv6_state - Is IPv6 support enabled.
- ipsec_state - Is IPSec support enabled.
- ipsec_track_tunnels - Should notify when IPSec tunnels are opened and closed.
- impending_power_loss_detection_method - The valid options are: BBU, None.
- async_min_rpo - Shortest ASync RPO allowed in the system.
- async_max_rpo - Longest ASync RPO allowed in the system.
- reserved_free_physical_capacity - Reserved physical space for OOPS recovery, in XAU units.

Example:

```
config_get
```

Output:

```
Name                          Value
----------------------- ----------------------------------
email_sender_address    support@example.com
email_reply_to_address  storage@example.com
dns_primary             192.0.2.1
dns_secondary
iscsi_name              iqn.2005-10.com.xivstorage:010140
system_name             IBM Storage System
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Allowed
Read-only users	Allowed

User Category	Permission
Technicians	Allowed

Return codes

- CONF_SERVER_UNREACHABLE
  The configuration server is unreachable.
- UNRECOGNIZED_CONFIG_PARAMETER
  Unrecognized configuration parameter: '*name*'.

  Troubleshooting: Use a valid configuration parameter in the command syntax. For the list of valid configuration parameters, see the CLI Reference Guide.

Setting configuration parameters

Use the **config_set** command to set configuration parameters.

```
config_set name=Name value=ParamValue
```

Parameters

Name	Type	Description	Mandatory
name	String	Name of the parameter to set.	Y
value	String	Value of the parameter.	Y

This command sets the values of configuration parameters.

The values of the following parameters can be set:

- dns_master - IP address of the master DNS server.
- dns_slave - IP address of the slave DNS server.
- email_sender_address - Email address used as the sender's address when sending email messages. Once set, this parameter cannot be set to null.
- email_reply_to_address - Reply-to address to be used when sending emails. This is useful for troubleshooting errors in email addresses.
- system_name - Name used as the sender's name when sending email messages.
- defaultuser - Default user to be used if no user is specified for the CLI. If null, a user must be specified.
- snmp_sysname - SNMP system name as shown in the SNMP MIB. (.1.3.6.1.2.1.1.5.0)
- snmp_location - SNMP location as shown in the SNMP MIB. (.1.3.6.1.2.1.1.6.0)
- snmp_contact - SNMP contact as shown in the SNMP MIB. (.1.3.6.1.2.1.1.4.0)
- email_subject_format - Controls the formatting of the email subject line. To insert the event's data, use the following tags: {severity}, {description}, or {system_name}. System default is "{severity}: {description}".
- ntp_server - IP address or DNS name of the NTP server.
- snmp_community - Community used for SNMP queries of the system.
- snmp_trap_community - Community used for SNMP traps sent by the system.
- snmp_trap_type - The valid options are: common, severity_specific.
- snmp_type - SNMP version the snmpd should work with. The valid options are: none, v2c, v3, both.

Example:

```
config_set name=dns_secondary value=192.0.2.2
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Allowed

Return codes

- UNRECOGNIZED_CONFIG_PARAMETER
  Unrecognized configuration parameter: '*name*'.

  Troubleshooting: Use a valid configuration parameter in the command syntax. For the list of valid configuration parameters, see the CLI Reference Guide.
- READ_ONLY_CONFIG_PARAMETER
  Configuration parameter: '*name*' is read-only.

Troubleshooting: Read-only parameters are not available for modifying.

- IPV4_NOT_CONFIGURED
  The IPv4 address is not configured on the management interface.

  Troubleshooting: Define an IPv4 address for management before disabling IPv6.

- RULE_WITH_SNMP_DEST_EXISTS
  Cannot set snmp_type to None. There is a rule that contains an SNMP destination.

Testing the DNS

Use the **dns_test** command to test the DNS (Domain Naming Service).

```
dns_test name=Name [ type=<A|AAAA> ]
```

Parameters

Name	Description	Mandatory	Default
name	Name of the host to be resolved.	Y	N/A
type	Type of query.	N	According to the DNS server type

This command attempts to translate the DNS name into an IP address. Translation is attempted through each of the defined DNS servers.

This command fails if no DNS servers are defined. A failure of the translation from a name to an IP address is not considered a failure of the command.

The result of each defined DNS server is displayed.

Field ID	Field output	Default position
name	Name	1
primary_ip	IP (Primary DNS)	2
secondary_ip	IP (Secondary DNS)	3

Example:

```
dns_test name=hermes.xiv
```

Output:

```
Name          IP (Primary DNS)   IP (Secondary DNS)
------------ ------------------ --------------------
hermes.xiv   203.0.113.1        Not Found
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Allowed

Return codes

- DNS_SERVER_NOT_DEFINED
  No DNS servers are defined.

Displaying help

Use the **help** command to display system help.

```
help [ category=Category | search=SearchString | command=CommandName ]
```

Parameters

Name	Type	Description	Mandatory
category	String	Category name.	N
search	String	Search string.	N
command	String	Command name.	N

This command displays the help as follows:

- No parameters - Lists all the commands with their short descriptions, grouped by categories.
- Category - Lists all the commands in the category, with their short descriptions.
- Search - Lists the short descriptions of all the commands in which the search string appears in their name or short description.
- Command - Displays the command name and short discrption.

Example:

```
help category=volume
```

Output:

```
Category  Name        Description

volume    vol_copy    Copies a source volume onto a target volume.

volume    vol_create  Creates a new volume.

volume    vol_delete  Deletes a volume

volume    vol_format  Formats a volume.

volume    vol_list    Lists all volumes, or a specific one.

volume    vol_lock    Locks a volume, so that it is read-only.

volume    vol_rename  Renames a volume

volume    vol_resize  Resizes a volume

volume    vol_unlock  Unlocks a volume, so that it is no longer read-only,
                      and can be written to.
```

Field ID	Field output	Default position
category	Category	1
name	Name	2
access_control	Access Control	N/A
syntax	Syntax	N/A
fields	Fields	N/A
description	Description	3
example	Example	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Allowed
Read-only users	Allowed
Technicians	Allowed

Printing the current maintenance urgency

Use the **maintenance_urgency_list** command to display the current maintenance urgency of the system.

```
maintenance_urgency_list
```

Example:

```
maintenance_urgency_list
```

Output:

```
maintenance_urgency = "NONE"
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed

User Category	Permission
Read-only users	Allowed
Technicians	Allowed

Shutting down the system

Use the **shutdown** command to shut down the system.

```
shutdown [ emergency=<yes|no> ] [ ignore_ha=<yes|no> ]
```

Parameters

Name	Type	Description	Mandatory	Default
emergency	Boolean	Instructs the system to shut down within a timeout even if some of the disks could not be saved, much like in an emergency shutdown performed when the system loses power.	N	no
ignore_ha	Boolean	Ignore activated HA objects.	N	no

The system stops serving hosts, de-stages all information to disks and then turns itself off. If the emergency parameter is defined, the system shuts down within the timeout period.

Attention:
USING THIS OPTION MAY CAUSE DATA LOSS.

Example:

```
shutdown -y
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_SHUT_DOWN
  Are you sure you want to shut down the machine and all its components?

Return codes

- COMMAND_IS_NOT_VALID_IN_CURRENT_SYSTEM_STATE
  The requested command cannot be invoked in the current system state.
- FIRMWARE_UPGRADE_IN_PROGRESS
  Firmware upgrade in progress.

  Troubleshooting: Contact IBM Support.
- FLASH_CCL_IN_PROGRESS
  The requested command cannot be invoked while Flash Enclosure CCL is in progress.

  Troubleshooting: Wait for Flash CCL to complete.
- SYSTEM_HAS_ACTIVE_MASTER_HA_RELATIONS
  The system owns volume(s) defined as primary in a HyperSwap relation. If you continue with shutdown without first handling those relations, the host may lose access to those volumes.

  Troubleshooting: It is recommended to run switch_roles before continuing. You may explicitly add ignore_ha=yes to force the operation.

Listing the operational state

Use the **state_list** command to display the current operational state of the system.

```
state_list
```

Example:

```
state_list
```

Output:

```
Category                 Value
------------------------ -----------------
data_protection_status   Fully Protected
data_reduction_state     Online
encryption               Enabled
safe_mode                no
shutdown_reason          No Shutdown
system_state             on
target_state             on
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Allowed
Read-only users	Allowed
Technicians	Allowed

Displaying system usage and data reduction statistics

Use the **system_usage_get** command to retrieve system-wide usage and data reduction related statistics.

```
system_usage_get
```

The command provides various information on system usage, and on data reduction, including:

- the sum of all user volume and snapshot sizes, excluding internal volumes (statistics and metadata)
- the sum of all logical block address (LBA) ranges written to the currently existing volumes and snapshots
- savings due to thin provisioning
- savings due to data reduction
- deduplication and compression factors.

The data retrieved by the command may vary due to currently running background processes.

Example:

```
system_usage_get
```

Output:

```
Volumes and Snapshots (GB)   Thin Provisioning Savings (%)   Total Written (GB)
--------------------------- ------------------------------ --------------------
3100                         90                             340

cont:
Total Written Pending Deletion (GB)   Data Reduction Savings (%)   Total Stored (GB)
------------------------------------ --------------------------- -------------------
0                                     87                          45

cont:
Deduplication Factor   Compression Factor
---------------------- --------------------
1.37                   5.63
```

Field ID	Field output	Default position
total_volumes_and_snapshots	Volumes and Snapshots (GB)	1
thin_provisioning_savings	Thin Provisioning Savings (%)	2
total_written	Total Written (GB)	3
total_written_pending_deletion	Total Written Pending Deletion (GB)	4
data_reduction_savings	Data Reduction Savings (%)	5
total_stored	Total Stored (GB)	6
deduplication_factor	Deduplication Factor	7
compression_factor	Compression Factor	8
data_only_deduplication_factor	Data Only Deduplication Factor	N/A
data_only_compression_factor	Data Only Compression Factor	N/A
data_only_reduction_factor	Data Only Reduction Factor	N/A

Field ID	Field output	Default position
deduplication_factor_full_accuracy	Deduplication Factor Full Accuracy	N/A
compression_factor_full_accuracy	Compression Factor Full Accuracy	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Displaying information about physical and effective capacity

Use the **system_capacity_list** command to display information about the system's effective and physical capacity.

```
system_capacity_list
```

The command output displays two lines of information: one for the effective capacity, and one for the physical capacity.

Effective capacity

Allocated effective capacity represents the sum of all virtual capacity provisioned to domains and to the pools in the default domain before any data reduction is applied.

There is a limit to the amount of effective capacity supported by each system. This limit is determined by the system configuration, namely, by the amount of grid controller resources, and it significantly exceeds the system physical capacity.

The command output for effective capacity includes the following field:

- Allocated (GB) — The sum of all effective capacity provisioned to domains, and to pools in the default domain. The value is represented in GB.

The rest of the output fields for effective capacity - Total, Allocated (%), Free (GB), and Free (%) - are currently not available (N/A).

Physical capacity

Physical capacity represents the amount of data that can be stored by the system after data reduction is applied. It is derived from the amount of flash storage media available in the system after taking into account the RAID protection scheme.

The command output for physical capacity includes the following fields:

- Total — The system's total physical capacity in GB.
- Allocated (GB) — The consumed physical capacity, represented in GB.
- Allocated (%) — The allocated capacity, represented as percentage of the total capacity.
- Free (GB) — The difference between the total and allocated capacity in GB.
- Free (%) — The free capacity, represented as percentage of the total capacity.

Example:

```
system_capacity_list
```

Output:

```
Type         Total (GB)   Allocated (GB) Allocated (%)  Free (GB)  Free (%)
------------ ------------ -------------- -------------- ---------- ---------
Effective    N/A          50640          N/A            N/A        N/A
Physical     1653         140            8              1513       92
```

Field ID	Field output	Default position
type	Type	1
total	Total (GB)	2
allocated	Allocated (GB)	3
allocated_percent	Allocated (%)	4
free	Free (GB)	5
free_percent	Free (%)	6
total_MiB	Total (MiB)	N/A
allocated_MiB	Allocated (MiB)	N/A
free_MiB	Free (MiB)	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Displaying information about effective capacity

Use the **system_effective_capacity_get** command to display additional information about the system's effective capacity.

```
system_effective_capacity_get
```

This command displays information about the system's effective capacity limit. This value depends on the current system configuration.

Example:

```
system_effective_capacity_get
```

Output:

```
Limit (GB)    Effective (GB)    Effective Factor
------------ ----------------- ------------------
1400071       N/A               N/A
```

Field ID	Field output	Default position
effective_capacity_limit	Limit (GB)	1
effective_capacity	Effective (GB)	2
effective_capacity_factor	Effective Factor	3
effective_capacity_max_limit	Max. Limit (GB)	N/A
effective_capacity_default_limit	Default Limit (GB)	N/A
effective_capacity_min_limit	Min. Limit (GB)	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Displaying system capacity thresholds

Use the **system_capacity_threshold_list** command to list the current system capacity thresholds.

```
system_capacity_threshold_list
```

A user can define up to 8 progressive thresholds that will trigger events about physical capacity consumption, as per the system_capacity_list command. An event of the configured severity is emitted once if the configured threshold value is exceeded, and an informational event is emitted when capacity recedes below the threshold minus the hysteresis.

For example, if a threshold is set at 85% and the hysteresis is set at 3%:

- a SYSTEM_CAPACITY_USAGE_INCREASED event is emitted when system capacity used in percent (per system_capacity_list) moves from a value below 85% to a value of 85% or more.
- the matching SYSTEM_CAPACITY_USAGE_DECREASED event is emitted when system capacity used moves from a value greater than 82% down to a value of 82% or less.

The output includes the following fields: This command displays the following information about the system capacity thresholds:

- The ordinal of the threshold (between 1 and 8)
- Threshold value in percent (between 10 and 99)
- Severity of the event
- Hysteresis value (between 1 and 10, same for all thresholds)
- Indication whether the threshold is enabled or not

Example:

```
system_capacity_threshold_list
```

Output:

```
Ordinal   Threshold (%)   Severity   Hysteresis (%)   Enabled
--------- --------------- ---------- ---------------- ---------
1         70              Warning    3                yes
2         75              Minor      3                yes
3         80              Minor      3                yes
4         85              Major      3                yes
5         90              Major      3                yes
6         95              Critical   3                yes
7         97              Critical   3                yes
8         99              Critical   3                yes
```

Field ID	Field output	Default position
ordinal	Ordinal	1
threshold	Threshold (%)	2
severity	Severity	3
hysteresis	Hysteresis (%)	4
enabled	Enabled	5

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Allowed

Changing a system capacity threshold

Use the **system_capacity_threshold_change** command to change a system capacity threshold.

```
system_capacity_threshold_change hysteresis=HysteresisValue |  < ordinal=Ordinal [ enabled=<yes|no> ]
 [ threshold=ThresholdValue ] [ severity=<WARNING|MINOR|MAJOR|CRITICAL> ] >
```

Parameters

Name	Type	Description	Mandatory	Default
hysteresis	Integer	The hysteresis value in per cent (same for all thresholds).	N	N/A
ordinal	Integer	The ordinal of the threshold.	N	N/A
enabled	Boolean	Enable or disable system capacity threshold.	N	yes
threshold	Integer	The new threshold value in percent, strictly monotonically increasing across thresholds.	N	No threshold
severity	N/A	The new severity value, strictly monotonically increasing across thresholds.	N	No severity

The user can define progressive thresholds that will trigger events about physical capacity consumption. An event of the configured severity is issued once if the configured threshold value is exceeded, and an informational all-clear event is issued, when capacity drops below the threshold minus the hysteresis.

The example below changes the hysteresis of the system capacity.

Example:

```
system_capacity_threshold_change hysteresis=3
```

The example below enables ordinal 1 and changes the threshold and severity parameters of the system capacity.

Example:

```
system_capacity_threshold_change ordinal=1 enabled=yes threshold=15 severity=WARNING
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- SYSTEM_CAPACITY_DUPLICATE_THRESHOLD
  The new threshold value is a duplicate of another threshold value.
- SYSTEM_CAPACITY_THRESHOLD_NOT_MONOTONIC
  The threshold value or severity must be monotone.

Resuming the system's normal operation

Use the **system_resume_normal_operation** command to move a system back to read-write state after it ran out of physical space.

```
system_resume_normal_operation
```

This command is used to move a system back to read-write state after it ran out of physical space.

Example:

```
system_resume_normal_operation
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_RESUME_NORMAL_OPERATION
  The system is still very close to full. Are you sure you want to resume normal operation?

Return codes

- SYSTEM_IS_STILL_OUT_OF_PHYSICAL_SPACE
  The system is still out of physical space, normal operation cannot be resumed.
- SYSTEM_ALREADY_OPERATING_NORMALLY
  The system is already operating normally.

Displaying the current time

Use the **time_list** command to display the current system time.

```
time_list
```

This command shows the current time, date and time zone.

Field ID	Field output	Default position
time	Time	1
date	Date	2
timezone	Time Zone	3
dst	Daylight Saving Time	4

Example:

```
time_list
```

Output:

```
Time       Date         Time Zone        Daylight Saving Time
---------- ------------ ---------------- ----------------------
10:09:47   2008-02-19   Asia/Jerusalem   no
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Setting the system's time

Use the **time_set** command to set the system's time in YYYY-MM-DD.HH:MM:SS format.

```
time_set time=Timestamp
```

Parameters

Name	Description	Mandatory
time	New current time.	Y

Example:

```
time_set time=2016-03-04.03:02:01
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Allowed

Return codes

- SYSTEM_TIME_NOT_CHANGED
  The system time has not changed.

  Troubleshooting: Please try again.
- FLASH_ENCLOSURE_TIME_UPDATE_FAILED
  Failed to update the flash enclosure time.

  Troubleshooting: Please try again.
- BAD_TIMESTAMP
  The timestamp cannot be deciphered.

Listing optional time zones

Use the **timezone_list** command to list all optional time zones.

```
timezone_list
```

Standard POSIX time zones are used. http://www.timeanddate.com/worldclock/ provides a full description of all time zones.

Example:

```
timezone_list
```

Output:

```
Timezone
---------------------------------
```

```
Africa/Abidjan
Africa/Accra

 ...

WET
Zulu
```

Field ID	Field output	Default position
name	Timezone	1

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Allowed
Read-only users	Allowed
Technicians	Allowed

Setting the time zone

Use the **timezone_set** command to set the time zone of the system.

```
timezone_set timezone=TimeZone
```

Parameters

Name	Type	Description	Mandatory
timezone	String	New time zone of the system.	Y

See Listing optional time zones for a complete list of optional time zones.

Standard POSIX time zones are used. http://www.timeanddate.com/worldclock/ provides a full description of all time zones.

Example:

```
timezone_set timezone=Etc/GMT+1
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Allowed

Return codes

- BAD_TIMEZONE_NAME
  Timezone is not recognized by the system.

Printing the current system version

Use the **version_get** command to print the current version of the system.

```
version_get
```

Field ID	Field output	Default position
system_version	Version	1

Example:

```
version_get
```

Output:

```
Version
10.2
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Allowed
Read-only users	Allowed
Technicians	Allowed

Displaying the values of VPD parameters

Use the **vpd_config_get** command to display the values of VPD parameters.

```
vpd_config_get  [ name=Name ]
```

Parameters

Name	Type	Description	Mandatory	Default
name	String	Name of the parameter to print.	N	All parameters.

Field ID	Field output	Default position
name	Name	1
value	Value	2

See Setting VPD parameters for a full list of available settings.

Example:

```
vpd_config_get name=site.city
```

Output:

```
Name         Value
-----------  --------
site.city    Gotham
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Allowed

Return codes

- CONF_SERVER_UNREACHABLE
  The configuration server is unreachable.
- UNRECOGNIZED_CONFIG_PARAMETER
  Unrecognized configuration parameter: '*name*'.

  Troubleshooting: Use a valid configuration parameter in the command syntax. For the list of valid configuration parameters, see the CLI Reference Guide.

Setting VPD parameters

Use the **vpd_config_set** command to set the values of VPD (Vital Product Data) parameters.

```
vpd_config_set name=Name value=ParamValue
```

Parameters

Name	Type	Description	Mandatory
name	String	Name of the parameter to set.	Y
value	String	Value of the parameter.	Y

This command sets the following values of VPD parameters, where only the **name** is mandatory.:

- `customer.name`
- `customer.primary_contact.calling_hours`
- `customer.primary_contact.email`
- `customer.primary_contact.mobile_phone`
- `customer.primary_contact.name`
- `customer.primary_contact.office_phone`
- `customer.primary_contact.time_zone`
- `customer.secondary_contact.calling_hours`
- `customer.secondary_contact.email`
- `customer.secondary_contact.mobile_phone`
- `customer.secondary_contact.name`
- `customer.secondary_contact.office_phone`
- `customer.secondary_contact.time_zone`
- `hardware_info.hw_cable_bundle`
- `hardware_info.hw_door`
- `hardware_info.hw_patch_panel`
- `hardware_info.hw_patch_panel_label`
- `hardware_info.hw_power_cable_config`
- `hardware_info.hw_rack_type`
- `hardware_info.hw_rps`
- `interface_config.model`
- `machine_model`
- `machine_type`
- `main_ibm_contact.calling_hours`
- `main_ibm_contact.email`
- `main_ibm_contact.mobile_phone`
- `main_ibm_contact.name`
- `main_ibm_contact.office_phone`
- `main_ibm_contact.time_zone`
- `non_mutable_vpd_info.original_flashed_version`
- `non_mutable_vpd_info.original_flashing_date`
- `disk_size`
- `remote_support.customer_contact.calling_hours`
- `remote_support.customer_contact.email`
- `remote_support.customer_contact.mobile_phone`
- `remote_support.customer_contact.name`
- `remote_support.customer_contact.office_phone`
- `remote_support.customer_contact.time_zone`
- `remote_support.modem_phone_number`
- `remote_support.primary_ibm_ip`
- `remote_support.secondary_ibm_ip`
- `remote_support.special_instructions`
- `remote_support.vpn_ip_1`
- `remote_support.vpn_ip_2`
- `site.building_location`
- `site.city`
  `site.country`
- `site.name`
- `site.postal_code`
- `site.state`
- `site.street_address`
- `system_info.sys_ec_level`
- `system_info.sys_hw_level`
- `system_info.machine_configuration_entry_point`
- `system_info.PID`

Example:

```
vpd_config_set name=customer.name value=JohnDoe
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed

User Category	Permission
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Allowed

Return codes

- READ_ONLY_CONFIG_PARAMETER
  Configuration parameter: '*name*' is read-only.

  Troubleshooting: Read-only parameters are not available for modifying.

- UNRECOGNIZED_CONFIG_PARAMETER
  Unrecognized configuration parameter: '*name*'.

  Troubleshooting: Use a valid configuration parameter in the command syntax. For the list of valid configuration parameters, see the CLI Reference Guide.

Displaying the system's MIB file

Use the **mib_get** command to display the system's MIB file.

```
mib_get
```

In configurations that use IBM® Netcool® Network Management for managing equipment, an enterprise (private) SNMP MIB from UC Davis is required. This MIB file can be downloaded from: http://www.net-snmp.org/docs/mibs/UCD-SNMP-MIB.txt.

After obtaining the **a9000.mib** file from the device with the **mib_get** command, note the following IMPORTS declaration line, which requires the parent UCD-SNMP-MIB:

Output:

```
IMPORTS
    MODULE-IDENTITY, OBJECT-TYPE,
        NOTIFICATION-TYPE,
        Gauge32, Integer32  FROM SNMPv2-SMI
                 ucdavis FROM UCD-SNMP-MIB
    OBJECT-GROUP, NOTIFICATION-GROUP,
        MODULE-COMPLIANCE    FROM SNMPv2-CONF
        sysName           FROM SNMPv2-MIB

    TEXTUAL-CONVENTION, DisplayString
                              FROM SNMPv2-TC;
```

When both MIB files (**a9000.mib** and **UCD-SNMP-MIB**) are imported into MIB Manager, the full OID path is properly defined by the combination of the declarations in each.

Field ID	Default position
line	1

Example:

```
mib_get
```

Output:

```
-------------------------------------------
-------------------------------------------
-------------------------------------------
-- -*- SNMP -*- mode for Emacs
XIV-MIB DEFINITIONS ::= BEGIN

IMPORTS
    MODULE-IDENTITY, OBJECT-TYPE,
        NOTIFICATION-TYPE,
        Gauge32, Integer32   FROM SNMPv2-SMI
                 ucdavis FROM UCD-SNMP-MIB
    OBJECT-GROUP, NOTIFICATION-GROUP,
        MODULE-COMPLIANCE    FROM SNMPv2-CONF

    TEXTUAL-CONVENTION, DisplayString
                              FROM SNMPv2-TC;
...
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed

User Category	Permission
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Allowed

Return codes

- CANNOT_READ_FROM_FILE
  Cannot read from file '*Filename*'.

  Troubleshooting: Contact IBM Support.

Retrieving the electronic license acceptance status

Use the **elicense_status_get** command to retrieve the electronic license acceptance status.

```
elicense_status_get
```

Example:

```
elicense_status_get
```

Output:

```
Status
----------
Accepted
```

Field ID	Field output	Default position
status	Status	1

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Retrieving a fragment of the electronic license file

Use the **elicense_blob_get** command to retrieve a fragment of the electronic license file.

```
elicense_blob_get beg=BeginIndex size=Number
```

Parameters

Name	Type	Description	Mandatory
beg	Positive integer	Beginning of the fragment in bytes.	Y
size	Positive integer	Length of the fragment in bytes. The maximum length allowed is 1000000.	Y

Example:

```
elicense_blob_get beg=0 size=20
```

Output:

```
<file_size value="1300473"/>
<fragment value="425a6839314159265359ba94ca1106dd587f84fe"/>
<fragment_size value="20"/>
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed

User Category	Permission
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Return codes

- CANNOT_READ_FROM_FILE
  Cannot read from file '*Filename*'.

  Troubleshooting: Contact IBM Support.

Accepting the electronic license agreement

Use the **elicense_accept** command to accept the electronic license agreement.

```
elicense_accept version=Version [ approver_name=UserName ]
```

Parameters

Name	Type	Description	Mandatory	Default
version	String	The electronic license version. For the instructions on retrieving the correct electronic license version, see below.	Y	N/A
approver_name	String	The approver's name.	N	none

To retrieve the correct electronic license version, proceed as follows:

1. Run the command **elicense_status_get -x**.
2. In the command output, find the string version value and copy its value.

Example:

```
elicense_accept version=xiv_license_v11.6.2_with_ela approver_name=johndoe
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- ELICENSE_INCOMPATIBLE_VERSION
  The accepted version of the Electronic license dose not match the current version

  Troubleshooting: Please retrieve the current electronic license version and accept it

- ELICENSE_ALREADY_ACCEPTED
  Electronic license already accepted

  Troubleshooting: You do not need to accept the electronic license

- ELICENSE_DISABLED
  Electronic license check is disabled

  Troubleshooting: You do not need to accept the electronic license

Enabling command auditing

Use the **audit_enable** command to enable CLI command auditing

```
audit_enable
```

This command is used by a security administrator to enable the auditing of user-entered CLI commands on an external auditing server. For this command to complete successfully, the current auditing state must be DISABLED (that is, the **audit_show** command returns a *no*), and at least one audit server must be

configured successfully by the **audit_config_set** command.

Note:
The **audit_enable** command has a default filter which filters out the following commands: **appadmin_capabilities_get**, **domain_global_list**, **state_list**, **metadata_list**, **snap_group_list**, **user_metadata_list**, and **version_get**. For more information about filters and filter changes, contact IBM Support.

Example:

```
audit_enable
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- AUDIT_ALREADY_ENABLED
  Command auditing is already enabled.
- AUDIT_NO_AUDIT_SERVER_DEFINED
  No audit logging server is configured.

Disabling command auditing

Use the **audit_disable** command to disable CLI command auditing.

```
audit_disable
```

This command disables command auditing, provided that auditing is currently enabled, that is the **audit_show** command returns a *yes*.

Example:

```
audit_disable -y
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- AUDIT_NOT_ENABLED
  Command auditing is not enabled.

Displaying the command audit state

Use the **audit_show** command to show the current state of CLI command auditing.

```
audit_show
```

Field ID	Field output	Default position

Field ID	Field output	Default position
audit_enabled	Auditing Enabled	1

Example:

```
audit_show
```

Output:

```
Auditing Enabled
------------------
yes
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Allowed
Read-only users	Allowed
Technicians	Allowed

Configuring audit servers

Use the **audit_config_set** command to configure CLI command auditing.

```
audit_config_set primary_server=Address [ primary_port=port ] [ secondary_server=Address ]
 [ secondary_port=port ] [ protocol=protocol ]
```

Parameters

Name	Type	Description	Mandatory	Default
primary_server	N/A	IP address of the primary auditing server.	Y	N/A
primary_port	Positive integer	IP port number of the primary auditing server.	N	Default for protocol
secondary_server	N/A	IP address of the secondary auditing server.	N	empty
secondary_port	Positive integer	IP port number of the secondary auditing server.	N	Default for protocol
protocol	Enumeration	Transport protocol. Only RFC-5424 Syslog over UDP is currently supported.	N	syslog

This command configures the primary and, optionally, the secondary auditing server for CLI command logging.

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Checking the command audit state

Use the **audit_config_get** command to show the current configuration of CLI command auditing.

```
audit_config_get
```

Field ID	Field output	Default position
primary_server	Primary Server	1
primary_port	Primary Port	2
secondary_server	Secondary Server	3
secondary_port	Secondary Port	4
audit_protocol	Protocol	5

Example:

```
audit_config_get
```

Output:

```
Primary Server   Primary Port   Secondary Server   Secondary Port   Protocol
---------------- -------------- ------------------ ---------------- ----------
192.0.2.1        514                               0                syslog
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Retrieving the list of Flash VDisks

Use the **flash_vdisk_list** command to retrieve the list of Flash VDisks

```
flash_vdisk_list
```

Example:

```
flash_vdisk_list
```

Output:

```
Name             Enclosure Id           Status   ID   Lun
---------------- ---------------------- -------- ---- ------------------
xiv_vdisk_2_0    1:Flash_Enclosure:2    OK       0    0000000000000000

Capacity(bytes)   Block Size
----------------- ------------
5717176090624     512
```

Field ID	Field output	Default position
name	Name	1
enclosure_id	Enclosure Id	2
status	Raid Status	3
vdisk_id	ID	4
lun	Lun	5
capacity_in_gb	Capacity(GB)	6
block_size	Block Size(bytes)	7
capacity_in_gib	Capacity(GiB)	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Enabling CIM service

Use the **cim_enable** command to enable the CIM service.

```
cim_enable
```

This command enables the CIM service and the associated SLP service. In order for this command to complete successfully, the current CIM service state must be DISABLED (that is, the **cim_show** command returns *no*).

Example:

```
cim_enable
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- CIM_ALREADY_ENABLED
  The CIM port is already enabled.

Disabling the CIM service

Use the **cim_disable** command to disable the CIM service.

```
cim_disable
```

This command disables the CIM service and the associated SLP service. In order for this command to complete successfully, the current CIM service state must be ENABLED (that is, the **cim_show** command returns *yes*).

Example:

```
cim_disable
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- CIM_NOT_ENABLED
  The CIM port is not enabled.

Displaying the CIM service state

Use the **cim_show** command to display the current state of CIM service.

```
cim_show
```

Field ID	Field output	Default position
cim_enabled	CIM Enabled	1

Example:

```
cim_show
```

Output:

```
CIM Enabled
-----------
yes
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Allowed
Read-only users	Allowed
Technicians	Allowed

Remote target connectivity commands

This section describes the command-line interface (CLI) for defining remote target connectivity.

- **Setting the threshold of a link disruption duration that triggers an event**
  Use the **target_change_connection_threshold** command to set the threshold of a link disruption that lasts more than a specified duration.
- **Updating the target's mirroring configuration**
  Use the **target_config_sync_rates** command to change the target's mirroring configuration.
- **Activating connectivity to a remote target**
  Use the **target_connectivity_activate** command to activate connectivity between a port on the local storage system and a port on a remote target.
- **Deactivating connectivity to a remote target**
  Use the **target_connectivity_deactivate** command to deactivate connectivity between a port on the local storage system and a port on a remote target.
- **Defining connectivity to a remote target**
  Use the **target_connectivity_define** command to define connectivity between a port on the local storage system and a port on a remote target.
- **Deleting connectivity to a remote target**
  Use the **target_connectivity_delete** command to delete connectivity between a port on the local storage system and a port on a remote target.
- **Listing target connectivity definitions**
  Use the **target_connectivity_list** command to list all the connectivity definitions of a remote target.
- **Defining a remote target**
  Use the **target_define** command to define a new remote target for remote mirroring or data migration.
- **Deleting a remote target**
  Use the **target_delete** command to delete the definition of the specified remote target.
- **Listing remote targets**
  Use the **target_list** command to list a specified remote target definition, or all target definitions.
- **Allowing remote mirroring access**
  Use the **target_mirroring_allow** command to allow remote mirroring operations initiated from a remote target.
- **Activating a port**
  Use the **target_port_activate** command to activate a port on a remote target.
- **Adding a new port to a remote target**
  Use the **target_port_add** command to add a port to a remote target.
- **Deactivating a port**
  Use the **target_port_deactivate** command to deactivate a port of a remote target.
- **Deleting a port from a remote system**
  Use the **target_port_delete** command to delete a port from the specified remote target.
- **Listing the ports of a remote target**
  Use the **target_port_list** command to list all ports of a target.
- **Renaming a remote target**
  Use the **target_rename** command to rename a remote target.
- **Updating the target configuration**
  Use the **target_update** command to update the target's configuration.
- **Adding a Quorum Witness to a target**
  Use the **target_add_quorum_witness** command to attach a Quorum Witness to a remote target.
- **Removing a Quorum Witness from a target**
  Use the **target_remove_quorum_witness** command to detach a Quorum Witness from a remote target.

Setting the threshold of a link disruption duration that triggers an event

Use the **target_change_connection_threshold** command to set the threshold of a link disruption that lasts more than a specified duration.

```
target_change_connection_threshold target=TargetName [ duration=duration ]
```

Parameters

Name	Type	Description	Mandatory	Default
duration	Integer	Duration for link down that will trigger an event, in seconds. Valid value is between 1 and 1000000 seconds.	N	30
target	Object name	The name of the target system for which the threshold is set.	Y	N/A

This command is used to set the duration of a link disruption that will trigger an event.

Example:

```
target_change_connection_threshold target="XIV MN00043" duration=25
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- TARGET_BAD_NAME
  The target name does not exist.
- TARGET_INVALID_CONNECTION_DURATION_THRESHOLD
  The target connection duration threshold should be in the [1,1000000] range.

Updating the target's mirroring configuration

Use the **target_config_sync_rates** command to change the target's mirroring configuration.

```
target_config_sync_rates target=TargetName [ max_initialization_rate=MaxInitializationRate ]

[ max_syncjob_rate=MaxSyncjobRate ] [ max_resync_rate=MaxResyncRate ]
```

Parameters

Name	Type	Description	Mandatory	Default
target	Object name	The updated target.	Y	N/A
max_initialization_rate	Positive integer	Specifies the maximum rate for initial synchronization. Cannot be larger than max_syncjob_rate.	N	Unchanged
max_syncjob_rate	Positive integer	Specifies the default maximum rate for sync job synchronization. Cannot be larger than max_resync_rate.	N	Unchanged
max_resync_rate	Positive integer	Specifies the maximum rate for re-synchronization	N	Unchanged

This command changes the system ID of the remote target. The synchronization rate units are MB per second. The default rates are: 100 MB/s for initialization rate, 300 MB/s for resync rate. The default system_id is the value that is set with the config_set command.

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- TARGET_BAD_NAME
  The target name does not exist.
- TARGET_ILLEGAL_RATE_VALUES
  The maximum initialization rate should be smaller than or equal to the maximum synchronization job rate. The maximum synchronization job rate should not be greater than the maximum resynchronization rate.

Activating connectivity to a remote target

Use the **target_connectivity_activate** command to activate connectivity between a port on the local storage system and a port on a remote target.

```
target_connectivity_activate target=TargetName
      < ipaddress=IPaddress local_ipinterface=IPInterface > |
```

```
<
      fcaddress=wwpn local_port=PortID >
```

Parameters

Name	Type	Description	Mandatory
target	Object name	Remote target of the connectivity definition.	Y
ipaddress	N/A	IP address of the port on the remote target (iSCSI targets only).	N
local_ipinterface	Object name	Local IP interface to be connected to the remote port (iSCSI only)	N
fcaddress	N/A	FC address of the port on the remote target (FC targets only).	N
local_port	N/A	Port identifier.	N

Each connectivity definition can be either active or inactive. The system does not use inactive connectivity definitions. Target connectivity is active by default.

This command has no effect if the connectivity is already active.

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- TARGET_BAD_NAME
  The target name does not exist.
- CONNECTIVITY_NOT_DEFINED
  No remote port is connected through this local port.
- COMPONENT_IS_NOT_AN_FC_PORT
  An FC port must be specified for the component.
- COMMAND_NOT_ALLOWED_ON_MANAGEMENT_OR_VPN_INTERFACE
  The operation is not allowed on the management or VPN IP Interface.
- IPINTERFACE_DOES_NOT_EXIST
  This IP interface name does not exist.
- TARGET_PORT_BAD_ADDRESS
  The remote port address is illegal or does not belong to the remote target.
- BAD_LOCAL_IP_PORT
  The ID of a local IP port must be specified.

Deactivating connectivity to a remote target

Use the **target_connectivity_deactivate** command to deactivate connectivity between a port on the local storage system and a port on a remote target.

```
target_connectivity_deactivate target=TargetName
< ipaddress=IPaddress local_ipinterface=IPInterface > |
<
fcaddress=wwpn  local_port=PortID > [ force_on_olvm_peer=<yes|no> ]
 [ force_on_ha_peer=<yes|no> ]
```

Parameters

Name	Type	Description	Mandatory	Default
target	Object name	Remote target of the connectivity definition.	Y	N/A
ipaddress	N/A	IP address of the port on the remote target (iSCSI targets only).	N	N/A
local_ipinterface	Object name	Local IP interface that is connected to the remote port (iSCSI only).	N	N/A
fcaddress	N/A	FC address of the port on the remote target (FC targets only).	N	N/A
local_port	N/A	Port identifier.	N	N/A
force_on_olvm_peer	Boolean	Informs the system whether the command should be applied on an olvm peer.	N	No
force_on_ha_peer	Boolean	Force the deactivation on a HyperSwap target.	N	No

This command deactivates connectivity.

Each connectivity definition can be either active or inactive. The system does not use inactive connectivity definitions. Target connectivity is active by default. Connectivity can be reactivated using Activating connectivity to a remote target.

This command has no effect if the connectivity is already deactivated.

Example:

```
target_connectivity_deactivate
target=Nextra2 local_module=101
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- TARGET_BAD_NAME
  The target name does not exist.
- COMMAND_NOT_ALLOWED_ON_MANAGEMENT_OR_VPN_INTERFACE
  The operation is not allowed on the management or VPN IP Interface.
- CONNECTIVITY_NOT_DEFINED
  No remote port is connected through this local port.
- COMPONENT_IS_NOT_AN_FC_PORT
  An FC port must be specified for the component.
- TARGET_PORT_BAD_ADDRESS
  The remote port address is illegal or does not belong to the remote target.
- BAD_LOCAL_IP_PORT
  The ID of a local IP port must be specified.
- IPINTERFACE_DOES_NOT_EXIST
  This IP interface name does not exist.
- TARGET_HAS_OLVM_RELATIONSHIP
  The target has an IBM Hyper-Scale Mobility relationship, and therefore cannot be deactivated or deleted.
- TARGET_HAS_HA_RELATIONSHIP
  The target has an IBM HyperSwap relationship, and therefore cannot be deactivated or deleted.

Defining connectivity to a remote target

Use the **target_connectivity_define** command to define connectivity between a port on the local storage system and a port on a remote target.

```
target_connectivity_define target=TargetName < ipaddress=IPaddress local_ipinterface=IPInterface > | < fcaddre
ss=wwpn local_port=PortID >
```

Parameters

Name	Type	Description	Mandatory
target	Object name	Remote target of the connectivity definition.	Y
ipaddress	N/A	IP address of the port on the remote target (iSCSI targets only).	N
local_ipinterface	Object name	Local IP interface to be connected to the remote port (iSCSI only).	N
fcaddress	N/A	FC address of the port on the remote target (FC targets only).	N
local_port	N/A	FC port (FC only).	N

Connectivity between a local and a target storage system is defined between a specific port on a local storage system and a port on the target storage system.

Each connectivity definition can be either active or inactive. The system does not use inactive connectivity definitions. Target connectivity is active by default. An option is provided to deactivate (target_connectivity_deactivate) and then reactivate (target_connectivity_activate) it, if required. Target connectivity can be deleted (Deleting connectivity to a remote target) and a list of target connectivity definitions (Listing target connectivity definitions) can be displayed.

If the port is used for target connectivity, then only one IP interface is supported. In this case, the other IP interfaces must be removed from the Ethernet port in order to define target connectivity.

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- TARGET_BAD_NAME
  The target name does not exist.
- CONN_EXISTS
  A remote port is already connected through this local port.
- MAX_CONNECTIONS_REACHED
  The maximum number of connections is already reached.
- MAX_ISCSI_CONNECTIONS_PER_MODULE_REACHED
  The maximum number of iSCSI connectivities for that module is already reached.
- COMPONENT_IS_NOT_AN_FC_PORT
  An FC port must be specified for the component.
- COMPONENT_IS_NOT_FC_INITIATOR_PORT
  An FC initiator port must be specified for the component.
- BAD_LOCAL_IP_PORT
  The ID of a local IP port must be specified.
- COMMAND_NOT_ALLOWED_ON_MANAGEMENT_OR_VPN_INTERFACE
  The operation is not allowed on the management or VPN IP Interface.
- MORE_THAN_ONE_IP_INTERFACE_DEFINED_ON_PORT
  There is more than one IP Interface defined on port.
- IPINTERFACE_DOES_NOT_EXIST
  This IP interface name does not exist.
- TARGET_PORT_BAD_ADDRESS
  The remote port address is illegal or does not belong to the remote target.

Deleting connectivity to a remote target

Use the **target_connectivity_delete** command to delete connectivity between a port on the local storage system and a port on a remote target.

```
target_connectivity_delete target=TargetName < ipaddress=IPaddress local_ipinterface=IPInterface > | < fcaddress=wwpn local_port=PortID > [ force_on_olvm_peer=<yes|no> ] [ force_on_ha_peer=<yes|no> ]
```

Parameters

Name	Type	Description	Mandatory	Default
target	Object name	Remote target of the connectivity definition.	Y	N/A
ipaddress	N/A	IP address of the port on the remote target (iSCSI targets only).	N	N/A
local_ipinterface	Object name	Local IP interface that is connected to the remote port (iSCSI only).	N	N/A
fcaddress	N/A	FC address of the port on the remote target (FC targets only).	N	N/A
local_port	N/A	Port number on the local module (FC only).	N	N/A
force_on_olvm_peer	Boolean	Informs the system whether the command should be applied on an IBM Hyper-Scale Mobility peer.	N	No
force_on_ha_peer	Boolean	Force the connectivity deletion on a HyperSwap target.	N	No

Only a previously defined connectivity definition can be deleted.

Example:

```
target_connectivity_delete target=XIV2 local_module=101
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- TARGET_BAD_NAME
  The target name does not exist.
- COMMAND_NOT_ALLOWED_ON_MANAGEMENT_OR_VPN_INTERFACE
  The operation is not allowed on the management or VPN IP Interface.
- CONNECTIVITY_NOT_DEFINED
  No remote port is connected through this local port.
- COMPONENT_IS_NOT_AN_FC_PORT
  An FC port must be specified for the component.
- TARGET_PORT_BAD_ADDRESS
  The remote port address is illegal or does not belong to the remote target.
- BAD_LOCAL_IP_PORT
  The ID of a local IP port must be specified.
- IPINTERFACE_DOES_NOT_EXIST
  This IP interface name does not exist.
- TARGET_HAS_OLVM_RELATIONSHIP
  The target has an IBM Hyper-Scale Mobility relationship, and therefore cannot be deactivated or deleted.
- TARGET_HAS_HA_RELATIONSHIP
  The target has an IBM HyperSwap relationship, and therefore cannot be deactivated or deleted.

Listing target connectivity definitions

Use the **target_connectivity_list** command to list all the connectivity definitions of a remote target.

```
target_connectivity_list [ target=TargetName ]  [ domain=DomainName ]
```

Parameters

Name	Type	Description	Mandatory	Default
target	Object name	Target name that is listed.	N	All targets
domain	Object name	The domain name.	N	All Domains

Field ID	Field output	Default position
target_name	Target Name	1
remote_port_address	Remote Port	2
local_fc_port	FC Port	3
local_ip_port	IP Interface	4
active	Active	5
up	Up	6

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Defining a remote target

Use the **target_define** command to define a new remote target for remote mirroring or data migration.

```
target_define target=TargetName  protocol=<FC|iSCSI> [ iscsi_name=iSCSIName ] [ xiv_features=<yes|no> ]
  [ system_id=SystemId ] [ domain=DomainList ] [ quorum_witness=QW_Name ]
 [ uses_512b_sectors=<yes|no> ]
```

Parameters

Name	Type	Description	Mandatory	Default
target	Object name	Local name of the remote target.	Y	N/A
protocol	Enumeration	FC (Fiber Channel) or iSCSI, depending on the communication protocol supported by the remote host.	Y	N/A
iscsi_name	iSCSI initiator name	iSCSI name of the remote target. This field is mandatory for iSCSI hosts.	N	N/A
system_id	String	ID of the remote system. Should be the same as the output of Displaying the values of configuration parameters of the *system_id* variable on the remote system.	N	N/A
xiv_features	Boolean	Defines the remote system as an XIV system. Non-XIV systems are used only for data migration.	N	Yes
domain	N/A	The cluster will be attached to the specified domains. To define more than one domain, separate them with a comma. To specify all existing domains, use "*".	N	none
quorum_witness	Object name	The name of the Quorum Witness that is associated with the target.	N	none
uses_512b_sectors	Boolean	Optimize the Asynchronous mirror data transfer for remote targets with 512B sector size.	N	No

This command defines the communication topology between a local storage system and a remote storage system to enable various features, such as remote mirroring. The local storage system can write to or read from the remote storage system, or allow the target storage system to write to or read from it.

The first step when defining a new target connectivity is to specify the name of the remote storage system and the protocol used to communicate with it. There are two possible protocols: Fiber Channel (FC) and iSCSI. Each remote target is available through only one of these protocols.

This step only defines the remote system object. No connectivity definitions are defined yet and no communications are performed yet.

Once you have defined a remote target, the only way to change its protocol type is to delete the remote target and define it again.

Example:

```
target_define target=Nextra2 protocol=FC
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_DEFINE_ANOTHER_TARGET_ON_SYSTEM
  Defining more than one target to the same remote system is not supported, and may compromise the data on the secondary system. Are you sure the remote system is not already defined as a target?

Return codes

- MAX_TARGETS_REACHED
  The maximum number of defined targets is already reached.
- TARGET_NAME_EXISTS
  The target name is already assigned to another target.
- TARGET_ISCSI_MUST_HAVE_A_NAME
  iSCSI Target must have an iscsi_name.
- ISCSI_NAME_NOT_ALLOWED_FOR_FC
  The FC Target does not have an iscsi_name.

- TARGET_BAD_SCSI_TYPE
  The target SCSI type does not exist.
- DOMAIN_DOESNT_EXIST
  The domain does not exist.
- QUORUM_WITNESS_BAD_NAME
  The Quorum Witness name does not exist.
- QUORUM_WITNESS_IS_NOT_ACTIVATED
  The Quorum Witness is not activated.
- QUORUM_WITNESS_CANNOT_BE_ADDED_TO_A_TARGET_OF_THIS_TYPE
  A Quorum Witness cannot be added to either iSCSI or non-Spectrum Accelerate target.

Deleting a remote target

Use the **target_delete** command to delete the definition of the specified remote target.

```
target_delete target=TargetName [ force_on_olvm_peer=<yes|no> ]
```

Parameters

Name	Type	Description	Mandatory	Default
target	Object name	Target that is deleted.	Y	N/A
force_on_olvm_peer	Boolean	Informs the system whether the command should be applied on an IBM Hyper-Scale Mobility peer.	N	No

A target that contains port definitions cannot be deleted. A target with remote mirroring or data migration definitions cannot be deleted.

Example:

```
target_delete target=Nextra2
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- TARGET_BAD_NAME
  The target name does not exist.
- TARGET_HAS_PORTS
  There are ports defined for this target.
- TARGET_HAS_ASSOCIATIONS
  There are remote volumes defined on this target.
- TARGET_HAS_OLVM_RELATIONSHIP
  The target has an IBM Hyper-Scale Mobility relationship, and therefore cannot be deactivated or deleted.
- TARGET_HAS_HA_RELATIONSHIP
  The target has an IBM HyperSwap relationship, and therefore cannot be deactivated or deleted.

Listing remote targets

Use the **target_list** command to list a specified remote target definition, or all target definitions.

```
target_list [ target=TargetName ] [ domain=DomainName ]
```

Parameters

Name	Type	Description	Mandatory	Default
target	Object name	Target name that is listed.	N	All targets
domain	Object name	The domain name.	N	All Domains

The following is listed for each target: port groups, ports, active/inactive status for each port, and the following mirroring-related values: max initialization rate, max resync rate, and max sync job rate.

Field ID	Field output	Default position
name	Name	1
scsi_type	SCSI Type	2
connected	Mirror Connectivity	3
ha_connected	HA Connectivity	4
max_initialization_rate	Max Initialization Rate	5
max_resync_rate	Max Resync Rate	6
max_syncjob_rate	Max Syncjob Rate	7
machine_serial_number	Target Serial Number	8
system_id	System ID	N/A
quorum_witness	Quorum Witness	9
xiv_target	XIV Target	N/A
iscsi_name	iSCSI Name	N/A
num_ports	Number of Ports	N/A
creator	Creator	N/A
connectivity_lost_event_threshold	Connection Threshold	N/A
peer_health	Peer Health	N/A
peer_health_reason	Peer Health Reason	N/A
peer_qw_configuration	Peer QW Configuration	N/A
coordinated_qw_lapse	Coordinated QW Lapse	N/A
arch	Remote Arch	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Allowing remote mirroring access

Use the **target_mirroring_allow** command to allow remote mirroring operations initiated from a remote target.

```
target_mirroring_allow target=TargetName
```

Parameters

Name	Type	Description	Mandatory
target	Object name	Remote target name.	Y

This command is performed on a local storage system in order to allow the target storage system to read, write, view, create volumes and define the existing volumes as slaves. This command is used when allowing remote mirroring operations. Otherwise, the target storage system cannot access the local storage system. This command also allows a remote target to read and write through the SCSI interface.

Once mirroring is allowed, this permission cannot be revoked.

This operation should also be run on the target storage system so that it gives permission to the local storage system to access it.

This step must be performed before mirroring is defined (mirror_create).

Example:

```
target_mirroring_allow target=Nextra2
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- TARGET_BAD_NAME
  The target name does not exist.
- TARGET_BAD_TYPE
  The target machine is not an XIV machine.

Activating a port

Use the **target_port_activate** command to activate a port on a remote target.

```
target_port_activate target=TargetName < ipaddress=IPaddress | fcaddress=wwpn >
```

Parameters

Name	Type	Description	Mandatory
target	Object name	Remote target of the port.	Y
ipaddress	N/A	IP address of the port on the remote target (iSCSI targets only).	N
fcaddress	N/A	FC address of the port on the remote target (FC targets only).	N

Each port in a remote system can be configured as either active or inactive. The system does not use inactive ports. After a port is defined, it is active by default. This command reactivates a port if it was de-activated (by using the target_port_deactivate command).

This command has no effect, if the port is already active.

Example:

```
target_port_activate target=Nextra2 fcaddress=10:00:00:17:38:27:ec:11
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- TARGET_PORT_BAD_ADDRESS
  The remote port address is illegal or does not belong to the remote target.
- TARGET_BAD_PORT_STATE
  The port is already in the requested activation state.
- TARGET_BAD_NAME
  The target name does not exist.

Adding a new port to a remote target

Use the **target_port_add** command to add a port to a remote target.

```
target_port_add target=TargetName < ipaddress=IPaddress | fcaddress=wwpn >
```

Parameters

Name	Type	Description	Mandatory
target	Object name	Remote target to which to add the port.	Y
ipaddress	N/A	IP address of the port on the remote target (for iSCSI type targets only).	N
fcaddress	N/A	FC address of the remote port (for FC type targets only).	N

This command adds a new port to a specified target. A port can be either FC or iSCSI, and its type must conform to the remote target's communication protocol type.

Specify the IP address or the FC address according to communication protocol of the target.

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- TARGET_BAD_NAME
  The target name does not exist.
- PORT_EXISTS
  The port is already defined.
- MAX_PORTS_REACHED
  The maximum number of ports defined in the system is already reached.
- TARGET_PORT_BAD_ADDRESS
  The remote port address is illegal or does not belong to the remote target.
- ISCSI_HOST_ILLEGAL_PORT_NAME
  The port name for iSCSI Host is illegal.

  Troubleshooting: Port names for iSCSI Hosts must contain only printable characters.
- HOST_PORT_EXISTS
  A host with this port ID is already defined.

Deactivating a port

Use the **target_port_deactivate** command to deactivate a port of a remote target.

```
target_port_deactivate target=TargetName < ipaddress=IPaddress | fcaddress=wwpn > [ force_on_olvm_peer=<yes|no
> ] [ force_on_ha_peer=<yes|no> ]
```

Parameters

Name	Type	Description	Mandatory	Default
target	Object name	The remote target that includes the port to be deactivated.	Y	N/A
ipaddress	N/A	IP address of the port on the remote target (iSCSI targets only).	N	N/A
fcaddress	N/A	FC address of the port on the remote target (FC targets only).	N	N/A
force_on_olvm_peer	Boolean	Informs the system whether the command should be applied on an OLVM peer.	N	No
force_on_ha_peer	Boolean	Force the deactivation on a HyperSwap target.	N	No

Each port in a remote system can be configured as either active or in-active. The system does not use an inactive port. After a port is defined, it is active by default. To re-activate a port, issue the target_port_activate command (see Activating a port).

Example:

```
target_port_deactivate target=XIV2 fcaddress=10:00:00:17:38:27:ec:11
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- TARGET_BAD_NAME
  The target name does not exist.
- TARGET_PORT_BAD_ADDRESS
  The remote port address is illegal or does not belong to the remote target.
- TARGET_BAD_PORT_STATE
  The port is already in the requested activation state.
- TARGET_HAS_OLVM_RELATIONSHIP
  The target has an IBM Hyper-Scale Mobility relationship, and therefore cannot be deactivated or deleted.
- TARGET_HAS_HA_RELATIONSHIP
  The target has an IBM HyperSwap relationship, and therefore cannot be deactivated or deleted.

Deleting a port from a remote system

Use the **target_port_delete** command to delete a port from the specified remote target.

```
target_port_delete target=TargetName < ipaddress=IPaddress | fcaddress=wwpn > [ force_on_ha_peer=<yes|no> ]
```

Parameters

Name	Type	Description	Mandatory	Default
target	Object name	Remote target from which the port is that is deleted.	Y	N/A
ipaddress	N/A	IP address of the port (for iSCSI targets only).	N	N/A
fcaddress	N/A	FC address of the remote port (for FC targets only).	N	N/A
force_on_ha_peer	Boolean	Force the deactivation on a HyperSwap target.	N	No

Example:

```
target_port_delete target=Nextra2 fcaddress=10:00:00:17:38:27:ec:11
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- TARGET_BAD_NAME
  The target name does not exist.
- TARGET_PORT_BAD_ADDRESS
  The remote port address is illegal or does not belong to the remote target.
- TARGET_PORT_HAS_CONNECTIVITY
  Connectivity to this port is already defined.
- TARGET_HAS_OLVM_RELATIONSHIP
  The target has an IBM Hyper-Scale Mobility relationship, and therefore cannot be deactivated or deleted.
- TARGET_HAS_HA_RELATIONSHIP

The target has an IBM HyperSwap relationship, and therefore cannot be deactivated or deleted.

Listing the ports of a remote target

Use the **target_port_list** command to list all ports of a target.

```
target_port_list [ target=TargetName ] [ domain=DomainName ]
```

Parameters

Name	Type	Description	Mandatory	Default
target	Object name	Target for which all ports should be listed.	N	All systems
domain	Object name	The domain name.	N	All Domains

Field ID	Field output	Default position
target_name	Target Name	1
scsi_type	Port Type	2
active	Active	3
fc_wwpn	WWPN	4
iscsi_ip_addr	iSCSI Address	5
iscsi_ip_port	iSCSI Port	6

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Renaming a remote target

Use the **target_rename** command to rename a remote target.

```
target_rename target=TargetName new_name=Name
```

Parameters

Name	Type	Description	Mandatory
target	Object name	The target to be renamed.	Y
new_name	Object name	New name of the target.	Y

Example:

```
target_rename target=Nextra2 new_name=Nextra-DRP
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- TARGET_BAD_NAME
  The target name does not exist.

- TARGET_NAME_EXISTS
  The target name is already assigned to another target.

Updating the target configuration

Use the **target_update** command to update the target's configuration.

```
target_update target=TargetName [ system_id=SystemId ] [ uses_512b_sectors=<yes|no> ]
```

Parameters

Name	Type	Description	Mandatory	Default
target	Object name	Target to be updated.	Y	N/A
system_id	String	ID of the remote system. Should be the same as the output of Displaying the values of configuration parameters of the *system_id* variable on the remote system.	N	none
uses_512b_sectors	Boolean	Optimize the Asynchronous mirror data transfer for remote targets with 512B sector size.	N	No

This command changes the system ID or other attributes of the remote target.

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- TARGET_BAD_NAME
  The target name does not exist.
- TARGET_HAS_QUORUM_WITNESS_UPDATE_NOT_ALLOWED
  Updating the target's system ID is not allowed when there is a Quorum Witness defined on the target.

Adding a Quorum Witness to a target

Use the **target_add_quorum_witness** command to attach a Quorum Witness to a remote target.

```
target_add_quorum_witness target=TargetName quorum_witness=QW_Name
```

Parameters

Name	Type	Description	Mandatory
target	Object name	The name of the target.	Y
quorum_witness	Object name	The name of the Quorum Witness that is associated with the target.	Y

Example:

```
target_add_quorum_witness target=t1 quorum=q1
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- TARGET_BAD_NAME
  The target name does not exist.
- QUORUM_WITNESS_BAD_NAME
  The Quorum Witness name does not exist.
- TARGET_HAS_A_QUORUM_WITNESS
  The target already has a Quorum Witness.
- QUORUM_WITNESS_IS_NOT_ACTIVATED
  The Quorum Witness is not activated.
- QUORUM_WITNESS_CANNOT_BE_ADDED_TO_A_TARGET_OF_THIS_TYPE
  A Quorum Witness cannot be added to either iSCSI or non-Spectrum Accelerate target.

Removing a Quorum Witness from a target

Use the **target_remove_quorum_witness** command to detach a Quorum Witness from a remote target.

```
target_remove_quorum_witness target=TargetName
```

Parameters

Name	Type	Description	Mandatory
target	Object name	The name of the target.	Y

Example:

```
target_remove_quorum_witness target=t1
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- TARGET_BAD_NAME
  The target name does not exist.
- TARGET_HAS_NO_QUORUM_WITNESS
  The local target does not have a Quorum Witness defined.
- TARGET_HAS_ENABLED_HA
  There are HyperSwap relations with enabled automatic failover configured with this target. Prior to changing the Quorum Witness, the automatic failover must be disabled.

Remote mirroring commands

This section describes the command-line interface (CLI) for remote mirroring.

Another command relevant to this topic is: Setting the threshold of a link disruption duration that triggers an event.

- **Canceling a snapshot mirror (ad hoc sync job)**
  Use the **mirror_cancel_snapshot** command to cancel all snapshot mirrors ('ad-hoc' sync jobs) of a specified master volume or a master consistency group, that have not run yet.
- **Creating a snapshot mirror (ad hoc sync job)**
  Use the **mirror_create_snapshot** command to create a snapshot mirror.
- **Activating mirroring**
  Use the **mirror_activate** command to activate mirroring for a defined mirror coupling.

- **Changing the RPO for local or remote system**
  Use the **mirror_change_rpo** command to change a local or remote RPO for a mirror relation.
- **Changing the designation of mirroring peers**
  Use the **mirror_change_designation** command to change the designation of mirroring peers: from primary to secondary, and vice versa.
- **Changing a mirroring schedule for remote slave peers**
  Use the **mirror_change_remote_schedule** command to change the replication schedule of a remote slave peer.
- **Changing the role of a mirrored volume**
  Use the **mirror_change_role** command to change the role of a local mirroring peer from Master to Slave or from Slave to Master.
- **Changing a mirroring schedule for local peers**
  Use the **mirror_change_schedule** command to change the replication schedule for peers on the local system.
- **Creating a mirroring definition**
  Use the **mirror_create** command to create a remote mirroring coupling.
- **Deactivating mirroring**
  Use the **mirror_deactivate** command to deactivate mirroring for a defined mirror coupling.
- **Deleting a remote mirroring definition**
  Use the **mirror_delete** command to delete a remote mirroring coupling definition.
- **Viewing the mirroring status**
  Use the **mirror_list** command to list the status and configuration of mirroring couplings.
- **Obtaining statistics on past sync jobs**
  Use the **mirror_statistics_get** command to present statistics that are automatically gathered by the system on past sync jobs per specified mirrored volume or consistency job.
- **Switching roles between master and slave**
  Use the **mirror_switch_roles** command to switch roles between master and slave volumes.
- **Retrieving RPO thresholds**
  Use the **rpo_thresholds_get** command to list system RPO-related thresholds, that, once crossed, trigger the creation of a corresponding event.
- **Setting an RPO threshold**
  Use the **rpo_thresholds_set** command to set system RPO-related thresholds, that, once crossed, trigger the creation of a corresponding event.
- **Changing the interval of a schedule**
  Use the **schedule_change** command to change the interval of a schedule.
- **Creating a schedule object**
  Use the **schedule_create** command to define a schedule for replication.
- **Triggering a schedule**
  Use the **schedule_create_tick** command to trigger a schedule-equivalent event for the couplings with the specified schedule.
- **Deleting a schedule object**
  Use the **schedule_delete** command to delete a schedule for replication.
- **Listing a schedule object**
  Use the **schedule_list** command to list the schedule properties for the specified coupling.
- **Renaming a schedule**
  Use the **schedule_rename** command to rename a schedule object.
- **Viewing sync job status**
  Use the **sync_job_list** command to list the statuses of queued and running sync jobs for asynchronous couplings

Canceling a snapshot mirror (ad hoc sync job)

Use the **mirror_cancel_snapshot** command to cancel all snapshot mirrors ('ad-hoc' sync jobs) of a specified master volume or a master consistency group, that have not run yet.

```
mirror_cancel_snapshot <vol=VolName | cg=cgName> [ target=TargetName ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	Name of the (local) master volume whose non-started snapshot mirrors should be cancelled.	N	N/A
cg	Object name	Name of the (local) master consistency group whose non-started snapshot mirrors should be cancelled.	N	N/A
target	Object name	Target mirror name. Mandatory if 2 mirrors are defined on the volume.	N	[none]

Only sync jobs that have not started are cancelled. The command does not delete the snapshots themselves.

Upon running the command:

- A warning message is presented to the user for confirmation.
- An event is generated.
- Non-started snapshot mirrors are cancelled.

The command fails under the following conditions:

- The command is issued on a slave volume or consistency group.

Access control

User Category	Permission	Condition
Storage administrator	Allowed	N/A

User Category	Permission	Condition
Storage integration administrator	Allowed	N/A
Application administrator	Conditionally Allowed	The volume is mapped to a host or a cluster associated with the user. If a snapshot overwrite is used, the target snapshot must be one created by a server administrator.
Security administrator	Disallowed	N/A
Read-only users	Disallowed	N/A
Technicians	Disallowed	N/A

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_CANCEL_SNAPSHOT_MIRRORS_FOR_THE_VOLUME
  Are you sure you want to delete snapshot mirrors for *Volume*?
- ARE_YOU_SURE_YOU_WANT_TO_CANCEL_SNAPSHOT_MIRRORS_FOR_THE_CONSISTENCY_GROUP
  Are you sure you want to delete snapshot mirrors for *Consistency Group*?

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.
- VOLUME_NO_MIRROR
  The local volume does not have remote mirroring definitions.
- CONS_GROUP_BAD_NAME
  The consistency group name does not exist.
- CONS_GROUP_NO_MIRROR
  The local consistency group does not have remote mirroring definitions.
- LOCAL_PEER_IS_NOT_MASTER
  The local peer is not primary.
- VOLUME_BELONGS_TO_MIRRORED_CONS_GROUP
  The volume mirror is part of a consistency group mirror.
- MIRROR_RETRY_OPERATION
  There is an operation in progress on this mirror.

  Troubleshooting: Retry the command in a few seconds.
- TARGET_BAD_NAME
  The target name does not exist.
- VOLUME_HAS_MULTIPLE_MIRRORS
  The volume has multiple mirrors. The operation is not allowed or a target must be specified.
- CONS_GROUP_HAS_MULTIPLE_MIRRORS
  The consistency group has multiple mirrors. The operation is not allowed, or a target must be specified.
- VOLUME_TARGET_MISMATCH
  The volume and target do not match.
- CONS_GROUP_BAD_TARGET
  The target name does not match the consistency group.
- MIRROR_IS_STANDBY
  The mirror is marked as Standby.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support

Creating a snapshot mirror (ad hoc sync job)

Use the **mirror_create_snapshot** command to create a snapshot mirror.

```
mirror_create_snapshot <vol=VolName | cg=cgName> [ target=TargetName ]
 < <name=Name [ delete_priority=del_value ]
 > | overwrite=Name> < < slave_name=SnapshotName [ slave_delete_priority=del_value ]
 > | slave_overwrite=SnapshotName>
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	The name of the volume to create a snapshot for.	N	N/A
cg	Object name	Local master consistency group name.	N	N/A
target	Object name	Target mirror name. Mandatory if 2 mirrors are defined on the volume.	N	[none]
name	Object name	The name of the new snapshot.	N	N/A
overwrite	Object name	The name of an existing snapshot that will be overwritten.	N	N/A
slave_overwrite	Object name	Name of existing snapshot on slave to overwrite.	N	N/A
delete_priority	Integer	The deletion priority of the volume's snapshot.	N	1
slave_name	Object name	The name of the new snapshot on the slave.	N	N/A
slave_delete_priority	Integer	The deletion priority of the slave volume's snapshot.	N	1

In synchronous replication, this command takes a snapshot of the source peer (master) and the target peer (slave) at exactly the same time.

In asynchronous replication, the command establishes a process that takes a point-in-time snapshot of the source peer (master) and synchronizes that point-in-time with the slave. The process sets a new sync job to copy the differences between that snapshot and the most recent snapshot that is guaranteed to be synchronized with the target peer.

Prerequisite (for both synchronous and asynchronous mirroring):

- The coupling has to be operational.

Multiple snapshot mirrors:

- Multiple snapshot mirrors can be issued; each mandates the creation of a corresponding sync job.
- Corresponding sync jobs are queued one after another.

Prioritization of sync jobs:

- The snapshot mirror delays the execution of an interval-based mirror if it is running upon arrival of a new interval.
- The snapshot mirror does not, however, cancel the creation of the interval-based sync job. The interval-based mirror will be calculated based on the differences between the most recent snapshot and the last snapshot mirror.

Precedence of the last snapshot mirror over the last replicated snapshot:

- The last_replicated snapshot of the master will be updated to reflect the completed snapshot mirror. Following the completion of the snapshot mirror, its snapshot is duplicated and the duplicate is named last_replicated (the previous last_replicated snapshot is deleted).

Canceling a snapshot mirror:

- The administrator has the ability to cancel snapshot mirrors that have not yet started.

Important:
The snapshots created concurrently on the master and slave are identical.

The snapshot mirror results with two last_replicated snapshots that are different and denoted "Master" and "Slave" accordingly:

- On the slave, a snapshot is taken and named last_replicated
- On the master, the pertinent snapshot that is mirrored onto the slave is also named last_replicated

The outcome for the synchronous mirroring:

- The master blocks host I/O for the duration of creating the snapshots
- The master completes synchronizing pending writes
- A snapshot of the master and slave is taken
- The master no longer blocks host I/O
- An event is generated

Using the overwrite and slave_overwrite parameters:

It is possible to overwrite an existing snapshot or snapshot group either on the Master, the Slave, or both.

To specify a local snapshot or snapshot group to be overwritten, use the overwrite parameter. Use the slave_overwrite parameter to specify a remote snapshot or snapshot group to be overwritten.

The overwrite and slave_overwrite parameters cause the current content of the volume or consistency group to be copied into an existing snapshot or snapshot groups (indicated as the parameter's argument). The overwritten snapshot or snapshots retain the same SCSI device WWN and the same mapping. As a result, the hosts maintain a continuous mapping of the snapshots, and a rescan or similar operation is not needed. The overwritten snapshot or snapshot group must be an existing snapshot or snapshot group of the respective volume or consistency group.

Access control

User Category	Permission	Condition
Storage administrator	Allowed	N/A
Storage integration administrator	Allowed	N/A
Application administrator	Conditionally Allowed	The volume is mapped to a host or a cluster associated with the user. If a snapshot overwrite is used, the target snapshot must be one created by a server administrator.
Security administrator	Disallowed	N/A
Read-only users	Disallowed	N/A
Technicians	Disallowed	N/A

Return codes

- CONS_GROUP_MISMATCH
  The snapshot group does not match the consistency group volumes.
- CONS_GROUP_EMPTY
  The operation is not allowed on an empty consistency group.
- CONS_GROUP_BAD_NAME
  The consistency group name does not exist.
- CONS_GROUP_NO_MIRROR
  The local consistency group does not have remote mirroring definitions.
- LOCAL_PEER_IS_NOT_MASTER
  The local peer is not primary.
- MIRROR_IS_NOT_SYNCHRONIZED
  The mirror is not synchronized.
- MIRROR_RETRY_OPERATION
  There is an operation in progress on this mirror.

  Troubleshooting: Retry the command in a few seconds.
- MIRROR_IS_NON_OPERATIONAL
  The mirror is non-operational.
- MAX_VOLUMES_REACHED
  The maximum allowed number of volumes is already reached.
- DOMAIN_MAX_VOLUMES_REACHED
  The domain exceeds the maximum allowed number of volumes.
- OPERATION_NOT_ALLOWED_ON_LOOPBACK
  The requested operation is not allowed on a loopback target.
- OVERWRITE_SNAPSHOT_BAD_NAME
  The snapshot name does not exist.
- OVERWRITE_SNAPSHOT_GROUP_DOES_NOT_BELONG_TO_GIVEN_GROUP
  The snapshot group belongs to another consistency group.
- POOL_SNAPSHOT_LIMIT_REACHED
  There is not enough space to create a snapshot.
- REMOTE_POOL_SNAPSHOT_LIMIT_REACHED
  There is not enough space on the remote target for creating a snapshot.
- REMOTE_MAX_VOLUMES_REACHED
  The maximum number of volumes on the remote machine is already reached.
- REMOTE_MAX_SNAPSHOTS_FOR_VOLUME_REACHED
  The maximum allowed number of snapshots per volume is already reached on a remote machine whose version is not 10.2.4.
- REMOTE_VOLUME_IS_MASTER
  A volume on the remote machine is already defined as primary.
- REMOTE_VOLUME_IS_SNAPSHOT
  The secondary volume is a snapshot.
- REMOTE_VOLUME_DATA_MIGRATION_UNSYNCHRONIZED
  Data Migration to the remote volume has not completed.
- REMOTE_SNAPSHOT_NAME_EXISTS
  The remote snapshot name already exists.
- REMOTE_SNAPSHOT_ILLEGAL_PRIORITY
  Illegal snapshot priority (remote); must be an integer between 1 and 4.
- REMOTE_SNAPSHOT_GROUP_NAME_EXISTS
  The remote snapshot group name already exists.
- REMOTE_SNAPSHOT_GROUP_ILLEGAL_PRIORITY
  Illegal snapshot group priority (remote); must be an integer between 1 and 4.
- REMOTE_SNAPSHOT_GROUP_BAD_PREFIX
  The remote snapshot group name has a reserved prefix.
- REMOTE_SNAPSHOT_BAD_PREFIX
  The remote snapshot name has a reserved prefix.
- REMOTE_CONS_GROUP_EMPTY

The operation is not allowed on an empty consistency group (remote).

- REMOTE_CONS_GROUP_MISMATCH
  The remote snapshot group does not match the consistency group volumes.
- SNAPSHOT_HAS_ACTIVE_SYNC_JOB
  The snapshot is currently the target of an active sync job.

  Troubleshooting: Please wait for the sync job to complete.
- SNAPSHOT_ILLEGAL_PRIORITY
  Illegal snapshot priority; must be an integer between 1 and 4.
- SNAPSHOT_IS_INTERNAL
  Internal snapshots cannot be mapped, modified or deleted.
- SNAPSHOT_GROUP_IS_INTERNAL
  Internal snapshots cannot be mapped, modified, or deleted.
- SNAPSHOT_GROUP_NAME_EXISTS
  The snapshot group name already exists.
- SNAPSHOT_GROUP_ILLEGAL_PRIORITY
  Illegal snapshot group priority; must be an integer between 1 and 4.
- SNAPSHOT_GROUP_BAD_NAME
  The snapshot group name does not exist.
- SNAPSHOT_GROUP_BAD_PREFIX
  The snapshot group name has a reserved prefix.
- SNAPSHOT_IS_PART_OF_SNAPSHOT_GROUP
  The snapshot is part of a snapshot group.
- SYNCHED_SNAPSHOTS_NOT_SUPPORTED_IN_TARGET
  The mirror's target does not support the synchronized snapshot capability.
- VOLUME_BAD_PREFIX
  The volume name has a reserved prefix.
- VOLUME_BELONGS_TO_MIRRORED_CONS_GROUP
  The volume mirror is part of a consistency group mirror.
- VOLUME_DATA_MIGRATION_UNSYNCHRONIZED
  Data Migration to this volume has not completed.
- VOLUME_EXISTS
  The volume name already exists.
- VOLUME_BAD_NAME
  The volume name does not exist.
- VOLUME_NO_MIRROR
  The local volume does not have remote mirroring definitions.
- VOLUME_IS_NOT_CONSISTENT_SLAVE
  The operation not allowed on an inconsistent secondary volume.
- VOLUME_IS_SNAPSHOT
  The operation is not permitted on snapshots.
- VOLUME_IS_OLVM_PROXY
  The volume is in an IBM Hyper-Scale Mobility Proxy phase.
- OPERATION_DENIED_OBJECT_MANAGED
  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.
- OPERATION_DENIED_REMOTE_OBJECT_MANAGED
  The remote object is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.
- CONS_GROUP_BAD_TARGET
  The target name does not match the consistency group.
- TARGET_BAD_NAME
  The target name does not exist.
- VOLUME_TARGET_MISMATCH
  The volume and target do not match.
- REMOTE_MIRROR_IS_STANDBY
  The remote mirror is marked as Standby.
- VOLUME_HAS_MULTIPLE_MIRRORS
  The volume has multiple mirrors. The operation is not allowed or a target must be specified.

- CONS_GROUP_HAS_MULTIPLE_MIRRORS
  The consistency group has multiple mirrors. The operation is not allowed, or a target must be specified.
- MIRROR_IS_STANDBY
  The mirror is marked as Standby.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support
- MAX_SNAPSHOTS_PER_VOLUME_REACHED
  The maximum allowed number of snapshots is already reached.
- REMOTE_MAX_SNAPSHOTS_PER_VOLUME_REACHED
  The maximum allowed number of snapshots is already reached on the remote ssytem.
- TARGET_SNAPSHOT_GROUP_BAD_NAME
  The target snapshot group name does not exist.
- REMOTE_DOMAIN_MAX_VOLUMES_REACHED
  The maximum number of volumes in the remote machine domain is already reached.
- REMOTE_CONS_GROUP_BAD_NAME
  The remote consistency group name does not exist.
- SNAPSHOT_CAN_NOT_BE_CREATED_REMOTE_CONS_GROUP_IO_IS_NOT_PAUSED
  The snapshot group will not be created since the remote consistency group is not in a stopped state.
- SNAPSHOT_CAN_NOT_BE_CREATED_REMOTE_CONS_GROUP_DEFINITION_CHANGED
  The snapshot group will not be created since the volumes in the remote consistency group have changed since the io_pause command was issued.
- REMOTE_OVERWRITE_SNAPSHOT_GROUP_DOES_NOT_BELONG_TO_GIVEN_GROUP
  The remote snapshot group belongs to another consistency group.
- REMOTE_SNAPSHOT_IS_PART_OF_SNAPSHOT_GROUP
  The subordinate snapshot is part of a snapshot group.
- REMOTE_SNAPSHOT_IS_INTERNAL
  Internal snapshot cannot be mapped, modified or deleted.
- REMOTE_SNAPSHOT_HAS_ACTIVE_SYNC_JOB
  The subordinate snapshot is currently the target of an active sync job.

  Troubleshooting: Please wait for the sync job to complete.
- REMOTE_OVERWRITE_SNAPSHOT_IS_MASTER_VOL
  The subordinate snapshot cannot be overwritten because it is a primary volume.
- REMOTE_OVERWRITE_SNAPSHOT_BAD_NAME
  The subordinate snapshot name does not exist.
- REMOTE_SNAPSHOT_OVERWRITE_MISMATCH
  The specified subordinate snapshot is not a snapshot of the specified volume.
- REMOTE_MULTISITE_IS_NOT_SMASTER
  The remote peer is not the SMaster.
- MULTISITE_BAD_GLOBAL_ID
  The Multi-site global ID does not exist.
- MULTISITE_SMASTER_INVALID_CONFIGURATION
  The Multi-site SMaster-Master relation configuration is invalid.

Activating mirroring

Use the **mirror_activate** command to activate mirroring for a defined mirror coupling.

```
mirror_activate < vol=VolName | cg=cgName > [ target=TargetName ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	Master volume.	N	N/A
cg	Object name	Master consistency group name or a list of master consistency groups.	N	N/A
target	Object name	Target mirror name. Mandatory if 2 mirrors are defined on the volume.	N	[none]

This command activates the coupling - either volumes or consistency groups - and switches it to the Active state.

Requirements for a successful command completion:

- The specified target must exist
- The specified target must be mirrored
- The specified target is a volume that does not belong to a consistency group, or is a consistency group
- The specified target is not a master
- The Standby state was explicitly set by issuing the mirror_deactivate command on the same peer

If the new activation state is the same as the existing state, nothing is done and a success code is returned.

The mirroring cannot be activated:

- If the time stamps of the last_replicated snapshots on the master and slave do not match.
- If the command is issued on a master that did not receive acknowledgment from the slave following the cg_add_vol or cg_remove_vol command (due to the command's timeout or to an unexpected failure), the command fails and the MIRROR_CONS_GROUP_MEMBERSHIP_MISMATCH code is returned. It means that the member lists of the mirror consistency group peers are not the same.
- If the command is issued on a master that did not receive acknowledgment from the slave following a vol_resize command (due to the command's timeout or to an unexpected failure), the command fails and the MIRROR_CONS_GROUP_MEMBERSHIP_MISMATCH code is returned. It means that the sizes of the mirror volume peers are not the same.

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.
- VOLUME_NO_MIRROR
  The local volume does not have remote mirroring definitions.
- CONS_GROUP_BAD_NAME
  The consistency group name does not exist.
- CONS_GROUP_NO_MIRROR
  The local consistency group does not have remote mirroring definitions.
- LOCAL_PEER_IS_NOT_MASTER
  The local peer is not primary.
- MIRROR_CONFIGURATION_ERROR
  The mirror's local configuration does not match its remote configuration.
- REMOTE_MAX_VOLUMES_REACHED
  The maximum number of volumes on the remote machine is already reached.
- SYNC_ALREADY_ACTIVE
  Synchronization is already active.
- VOLUME_BELONGS_TO_MIRRORED_CONS_GROUP
  The volume mirror is part of a consistency group mirror.
- DOMAIN_MAX_VOLUMES_REACHED
  The domain exceeds the maximum allowed number of volumes.
- MIRROR_CAN_NOT_BE_ACTIVATED
  Mirroring cannot be activated.
- MIRROR_CONS_GROUP_MEMBERSHIP_MISMATCH
  The mirrored consistency group contains different volumes on the primary and secondary machines. This problem occurs whenever the cg_add_vol or cg_remove_vol commands were previously issued, and the primary machine did not receive an acknowledgment from the secondary machine until the command timed out, or due to any other unexpected failure.
- MIRROR_SIZE_MISMATCH
  The secondary and primary volume sizes are different.
- MIRROR_RETRY_OPERATION
  There is an operation in progress on this mirror.

  Troubleshooting: Retry the command in a few seconds.
- VOLUME_HAS_MULTIPLE_MIRRORS
  The volume has multiple mirrors. The operation is not allowed or a target must be specified.

- CONS_GROUP_HAS_MULTIPLE_MIRRORS
  The consistency group has multiple mirrors. The operation is not allowed, or a target must be specified.
- TARGET_BAD_NAME
  The target name does not exist.
- VOLUME_TARGET_MISMATCH
  The volume and target do not match.
- CONS_GROUP_BAD_TARGET
  The target name does not match the consistency group.
- VOLUME_TOO_MANY_ACTIVE_MIRRORS
  This command cannot be issued if more than one mirror is active on the volume.
- REMOTE_MIRROR_IS_STANDBY
  The remote mirror is marked as Standby.
- REMOTE_DOMAIN_MAX_VOLUMES_REACHED
  The maximum number of volumes in the remote machine domain is already reached.
- MIRROR_IS_STANDBY
  The mirror is marked as Standby.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support
- REMOTE_DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier of the remote system is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support
- REMOTE_SYSTEM_OUT_OF_PHYSICAL_SPACE
  The operation not allowed while the remote system is out of physical space.
- CONS_GROUP_MEMBER_VOL_IS_MISSING_A_RELATION
  A volume in the consistency group is missing the volume-level Multi-site/mirror/HyperSwap relation.
- REMOTE_MULTISITE_IS_NOT_SMASTER
  The remote peer is not the SMaster.
- MULTISITE_SMASTER_INVALID_CONFIGURATION
  The Multi-site SMaster-Master relation configuration is invalid.
- MULTISITE_BAD_GLOBAL_ID
  The Multi-site global ID does not exist.

Changing the RPO for local or remote system

Use the **mirror_change_rpo** command to change a local or remote RPO for a mirror relation.

```
mirror_change_rpo  <vol=VolName | cg=cgName> [ target=TargetName ] [ rpo=rpo ] [ remote_rpo=rpo ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	Local volume name. Must be specified if the command is applied to a volume.	N	N/A
cg	Object name	Consistency group name on the local system.	N	N/A
target	Object name	Target name of the mirror, mandatory if there are 2 mirrors defined on the volume.	N	[none]
remote_rpo	Integer	RPO on a remote system.	N	[Unchanged]
rpo	Integer	RPO on the local system	N	[Unchanged]

- The command must be run on the master.
- The RPO must be greater than the interval.
- The link has to be up.

Example:

```
mirror_change_rpo vol=volname rpo=100
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.
- CONS_GROUP_NO_MIRROR
  The local consistency group does not have remote mirroring definitions.
- ASYNC_MIRROR_REMOTE_RPO_TOO_SHORT
  The specified remote RPO is too short.
- ASYNC_MIRROR_RPO_TOO_LONG
  The specified RPO is too long.
- TARGET_NOT_CONNECTED
  There is currently no connection to the target system.
- VOLUME_NO_MIRROR
  The local volume does not have remote mirroring definitions.
- CONS_GROUP_BAD_NAME
  The consistency group name does not exist.
- ASYNC_MIRROR_RPO_TOO_SHORT
  The specified RPO is too short.
- VOLUME_BELONGS_TO_MIRRORED_CONS_GROUP
  The volume mirror is part of a consistency group mirror.
- INTERVAL_SHOULD_BE_SHORTER_THAN_RPO
  The schedule interval must be shorter than the RPO.
- ASYNC_MIRROR_REMOTE_RPO_TOO_LONG
  The specified remote RPO is too long.
- LOCAL_IS_SLAVE
  The local mirror peer is not primary.
- SYNC_MIRROR_HAS_NO_RPO
  The synchronous mirror does not have an RPO.
- TARGET_BAD_NAME
  The target name does not exist.
- VOLUME_HAS_MULTIPLE_MIRRORS
  The volume has multiple mirrors. The operation is not allowed or a target must be specified.
- CONS_GROUP_HAS_MULTIPLE_MIRRORS
  The consistency group has multiple mirrors. The operation is not allowed, or a target must be specified.
- VOLUME_TARGET_MISMATCH
  The volume and target do not match.
- CONS_GROUP_BAD_TARGET
  The target name does not match the consistency group.
- MIRROR_IS_STANDBY
  The mirror is marked as Standby.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support

Changing the designation of mirroring peers

Use the **mirror_change_designation** command to change the designation of mirroring peers: from primary to secondary, and vice versa.

```
mirror_change_designation < vol=VolName | cg=cgName > [ target=TargetName ]
 [ new_designation=<Primary|Secondary|None> ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	Master volume name.	N	N/A
cg	Object name	Master consistency group name.	N	N/A
target	Object name	Target mirror name. Mandatory if 2 mirrors are defined on the volume.	N	[none]
new_designation	Enumeration	The new designation of the peer If not specified, the command swaps the designation of the primary and secondary peer.	N	none

The command is issued on the master peer and affects both peers. The coupling has to be operational.

The designation change implied by this command reflects a decision to reset the designation of the mirroring peers, in contrast with the operational role, which is denoted by the master/dlave title.

There is no obligation to issue the command with a specification of the new designation. If the new designation is not specified, the command swaps the designations of both peers from their current value. The primary changes to secondary, and the secondary - to primary.

Example:

```
 mirror_change_designation cg=reggie13_cg new_designation=Secondary
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.
- VOLUME_NO_MIRROR
  The local volume does not have remote mirroring definitions.
- CONS_GROUP_BAD_NAME
  The consistency group name does not exist.
- CONS_GROUP_NO_MIRROR
  The local consistency group does not have remote mirroring definitions.
- LOCAL_PEER_IS_NOT_MASTER
  The local peer is not primary.
- MIRROR_DESIGNATION_NOT_SUPPORTED_BY_TARGET
  The mirror's target does not support mirror role designation.
- MIRROR_IS_NON_OPERATIONAL
  The mirror is non-operational.
- VOLUME_BELONGS_TO_MIRRORED_CONS_GROUP
  The volume mirror is part of a consistency group mirror.
- TARGET_BAD_NAME
  The target name does not exist.
- VOLUME_HAS_MULTIPLE_MIRRORS
  The volume has multiple mirrors. The operation is not allowed or a target must be specified.
- CONS_GROUP_HAS_MULTIPLE_MIRRORS
  The consistency group has multiple mirrors. The operation is not allowed, or a target must be specified.
- VOLUME_TARGET_MISMATCH
  The volume and target do not match.
- CONS_GROUP_BAD_TARGET

The target name does not match the consistency group.

- MIRROR_IS_STANDBY
  The mirror is marked as Standby.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support

Changing a mirroring schedule for remote slave peers

Use the **mirror_change_remote_schedule** command to change the replication schedule of a remote slave peer.

```
mirror_change_remote_schedule < vol=VolName | cg=cgName > [ target=TargetName ] remote_schedule=Schedule
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	Local master volume name.	N	N/A
cg	Object name	Local master consistency group name.	N	N/A
target	Object name	Target mirror name. Mandatory if 2 mirrors are defined on the volume.	N	[none]
remote_schedule	Object name	A reference to a remote schedule that should be set for the remote slave peer, which corresponds with the master specified in the command.	Y	N/A

This command changes the replication schedule of an asynchronous coupling in order to make it effective after the role of a specified remote slave peer is changed to master.

Prerequisites:

- The coupling must be `ASYNC_INTERVAL`.

Following the command execution:

- The system displays a warning
- If the command is approved, it is executed
- An event is generated
- New sync jobs are generated according to the updated schedule
- Existing sync jobs are not affected (that is, they run according to the previous schedule)

Requirements for a successful command completion:

- The specified target exists
- The specified target is mirrored
- The specified target is not a volume that belongs to a mirrored consistency group
- The specified target is of sync type `ASYNC_INTERVAL`
- The specified target is a master
- The link is up

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.
- VOLUME_NO_MIRROR
  The local volume does not have remote mirroring definitions.
- CONS_GROUP_BAD_NAME
  The consistency group name does not exist.
- CONS_GROUP_NO_MIRROR
  The local consistency group does not have remote mirroring definitions.
- REMOTE_VOLUME_IS_MASTER
  A volume on the remote machine is already defined as primary.

- REMOTE_CONS_GROUP_IS_MASTER
  The remote consistency group is defined as primary.
- SCHEDULE_DOES_NOT_EXIST
  The specified schedule does not exist.
- SYNC_MIRROR_DOES_NOT_USE_SCHEDULE
  A synchronous mirror definition does not require a schedule object.
- TARGET_NOT_CONNECTED
  There is currently no connection to the target system.
- VOLUME_BELONGS_TO_MIRRORED_CONS_GROUP
  The volume mirror is part of a consistency group mirror.
- INTERVAL_SHOULD_BE_SHORTER_THAN_RPO
  The schedule interval must be shorter than the RPO.
- TARGET_BAD_NAME
  The target name does not exist.
- VOLUME_HAS_MULTIPLE_MIRRORS
  The volume has multiple mirrors. The operation is not allowed or a target must be specified.
- CONS_GROUP_HAS_MULTIPLE_MIRRORS
  The consistency group has multiple mirrors. The operation is not allowed, or a target must be specified.
- VOLUME_TARGET_MISMATCH
  The volume and target do not match.
- CONS_GROUP_BAD_TARGET
  The target name does not match the consistency group.
- DOMAIN_HAS_NO_ACCESS_TO_SCHEDULE
  The domain has no access to the schedule.
- MIRROR_IS_STANDBY
  The mirror is marked as Standby.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support

Changing the role of a mirrored volume

Use the **mirror_change_role** command to change the role of a local mirroring peer from Master to Slave or from Slave to Master.

```
mirror_change_role <vol=VolName | cg=cgName> [ target=TargetName ] [ new_role=<Master|Slave|None> ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	Local volume name. Must be specified if the command is applied to a volume.	N	N/A
cg	Object name	Consistency group name. Must be specified if the command is applied to a consistency group.	N	N/A
target	Object name	Target mirror name. Mandatory if 2 mirrors are defined on the volume.	N	[none]
new_role	Enumeration	Role name of the peer. If not specified, the command swaps peer roles between Master and Slave.	N	none

This command changes the role of the local peer from Master to Slave or from Slave to Master when the coupling is non-operational. It is assumed that the command will be issued on both peers of the coupling before the coupling becomes operational again, so that upon reconnection there still will be one Master and one Slave.

For a successful command completion:

- Do not issue the command on a peer whose status is Initializing.
- Do not issue the command in the Change Tracking state.
- Do not issue the command on a volume that belongs to a mirrored consistency group, otherwise the command will return an error and fail.

Changing the roles in synchronous mirroring:

When applied on a Master, the Master becomes a Slave, ceases serving host requests, and is set to accept replication from the other peer as a Slave.

When applied on a Slave, the Slave becomes a Master, starts accepting requests from hosts, and upon explicit activation starts replicating to the other peer (the original Master).

If the synchronous mirroring is interrupted in the middle of the re-synchronization process, the Slave volume may very probably be inconsistent. The last consistent image of the Slave volume is preserved in the last_consistent snapshot (LCS), which is automatically created immediately before the re-synchronization starts. If the LCS exists, the command emits a warning: Are you sure you want the mirror/HyperSwap local peer to become primary? The local peer has a last-consistent snapshot. In this case, the administrator must choose whether to use the existing contents of the previous Slave volume, which may be inconsistent, or revert the previous Slave volume to its last_consistent snapshot before issuing the mirror_change_role command.

Changing the roles in asynchronous mirroring:

When successfully applied on a Master, the Master is reverted to the image recorded on the last_replicated snapshot of the mirror, and ceases accepting host requests.

When applied on a Slave:

- A warning is displayed.
- An event is generated.
- The new Master ceases accepting replication requests from the previous Master, and reverts to the last_replicated snapshot.
- The new Master starts accepting host requests.
- The process completion is recorded in the log.
- Explicit activation of mirroring is required.

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- SOME_DATA_WILL_BE_LOST_ARE_YOU_SURE
  Are you sure you want the mirror/HyperSwap local peer to become secondary and lose the data that was not replicated?
- ARE_YOU_SURE_YOU_WANT_TO_CHANGE_A_PEER_WITH_LCS_TO_MASTER
  Are you sure you want the mirror/HyperSwap local peer to become primary? The local peer has a last-consistent snapshot.

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.
- VOLUME_NO_MIRROR
  The local volume does not have remote mirroring definitions.
- CONS_GROUP_BAD_NAME
  The consistency group name does not exist.
- CONS_GROUP_NO_MIRROR
  The local consistency group does not have remote mirroring definitions.
- MIRROR_IS_INITIAL
  The operation is not permitted during the Initialization phase.
- MIRROR_IS_ACTIVE
  Remote mirroring is currently active.
- VOLUME_HAS_DATA_MIGRATION
  Data Migration is defined for this volume.
- VOLUME_BELONGS_TO_MIRRORED_CONS_GROUP
  The volume mirror is part of a consistency group mirror.
- MIRROR_RETRY_OPERATION
  There is an operation in progress on this mirror.

  Troubleshooting: Retry the command in a few seconds.
- MIRROR_HAS_NO_SYNCHED_SNAPSHOT
  The mirror does not have a synchronized snapshot.
- MASTER_CANNOT_BE_DEMOTED
  The primary volume cannot be demoted to secondary. Peer status mismatch.
- VOLUME_HAS_MULTIPLE_RELATIONS
  The volume has multiple relations. The operation is not allowed or a target must be specified.

- CONS_GROUP_HAS_MULTIPLE_RELATIONS
  The consistency group has multiple relations. The operation is not allowed or a target must be specified.
- TARGET_BAD_NAME
  The target name does not exist.
- VOLUME_TARGET_MISMATCH
  The volume and target do not match.
- CONS_GROUP_BAD_TARGET
  The target name does not match the consistency group.
- MIRROR_PART_OF_MULTISITE
  The remote mirror is part of a Multi-site.
- MIRROR_IS_STANDBY
  The mirror is marked as Standby.
- COMMAND_NOT_SUPPORTED_FOR_OLVM_VOLUMES
  This command is not supported for IBM Hyper-Scale Mobility volumes.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support
- SYSTEM_OUT_OF_PHYSICAL_SPACE
  The operation not allowed while the system is out of physical space.

Changing a mirroring schedule for local peers

Use the **mirror_change_schedule** command to change the replication schedule for peers on the local system.

```
mirror_change_schedule < vol=VolName | cg=cgName > [ target=TargetName ] schedule=Schedule
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	Volume name on the local system.	N	N/A
cg	Object name	Consistency group name on the local system.	N	N/A
target	Object name	Target mirror name. Mandatory if 2 mirrors are defined on the volume.	N	[none]
schedule	Object name	A reference to a mirroring schedule	Y	N/A

This command changes the replication schedule for a peer on the local system. The new scheduling will become effective only if the peer is set as master.

Prerequisites:

- The coupling must be `ASYNC_INTERVAL`.
- The schedule's interval has to be shorter than the corresponding mirror's RPO.

The command fails under the following conditions:

- The specified target does not exist
- The specified target is non-mirrored
- The specified target is a volume that belongs to a mirrored Consistency Group
- The specified target synchronization type is not `ASYNC_INTERVAL`

Setting a scheduling reference:

- An event is generated
- New sync jobs will be generated according to updated schedule. A running sync job is unaffected.

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- VOLUME_BAD_NAME

The volume name does not exist.

- CONS_GROUP_BAD_NAME
  The consistency group name does not exist.
- VOLUME_NO_MIRROR
  The local volume does not have remote mirroring definitions.
- CONS_GROUP_NO_MIRROR
  The local consistency group does not have remote mirroring definitions.
- SCHEDULE_DOES_NOT_EXIST
  The specified schedule does not exist.
- SYNC_MIRROR_DOES_NOT_USE_SCHEDULE
  A synchronous mirror definition does not require a schedule object.
- VOLUME_BELONGS_TO_MIRRORED_CONS_GROUP
  The volume mirror is part of a consistency group mirror.
- INTERVAL_SHOULD_BE_SHORTER_THAN_RPO
  The schedule interval must be shorter than the RPO.
- ILLEGAL_INTERVAL
  The specified interval value is not supported.
- TARGET_BAD_NAME
  The target name does not exist.
- VOLUME_HAS_MULTIPLE_MIRRORS
  The volume has multiple mirrors. The operation is not allowed or a target must be specified.
- CONS_GROUP_HAS_MULTIPLE_MIRRORS
  The consistency group has multiple mirrors. The operation is not allowed, or a target must be specified.
- VOLUME_TARGET_MISMATCH
  The volume and target do not match.
- CONS_GROUP_BAD_TARGET
  The target name does not match the consistency group.
- MIRROR_IS_STANDBY
  The mirror is marked as Standby.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support

Creating a mirroring definition

Use the **mirror_create** command to create a remote mirroring coupling.

```
mirror_create < vol=VolName slave_vol=SlaveVolumeName [ create_slave=<yes|no> [ remote_pool=RemotePoolName ] ]
 [ init_type=<online|offline> ]
 > | <cg=cgName slave_cg=SlaveCgName> [ type=<SYNC_BEST_EFFORT|ASYNC_INTERVAL> ]
 target=TargetName [ rpo=rpo [ remote_rpo=rpo ] schedule=Schedule remote_schedule=Schedule ]
 [ part_of_multisite=<yes|no> ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	Local volume to be mirrored (the master).	N	N/A
slave_vol	Object name	The name of the slave volume on the remote storage system.	N	N/A
create_slave	Boolean	Determines whether to create a new slave volume or to use an existing one.	N	no
remote_pool	Object name	The storage pool on the remote system. Relevant only if creating a slave.	N	N/A
cg	Object name	Local consistency group to be mirrored (the master).	N	N/A
slave_cg	Object name	The name of the slave consistency group on the remote storage system.	N	N/A
type	Enumeration	The name of the replication type	N	SYNC_BEST_ EFFORT
target	Object name	Remote target to contain the slave volume.	Y	N/A
rpo	Positive integer	A mirror recovery point objective value for the master. Ranges from 30 to 86400 seconds (that is, up to 24 hours) Is applicable and mandatory for asynchronous mirroring only.	N	[None]
remote_rpo	Positive integer	Mirror recovery point objective value for a remote peer that becomes master Is applicable and mandatory for asynchronous mirroring only.	N	[Master RPO]

Name	Type	Description	Mandatory	Default
schedule	Object name	A reference to a schedule object Is applicable and mandatory for asynchronous mirroring only.	N	[None]
remote_schedule	Object name	A reference to a schedule object on the remote machine. Is applicable and mandatory for asynchronous mirroring only.	N	[None]
init_type	Enumeration	Specifies the method requested to initialize the slave mirror.	N	[none]
part_of_multisite	Boolean	Marks the mirror as part of a Multi-site relation.	N	no

Mirroring is the process of ensuring that both peers contain identical data at all times. This command defines a new mirroring coupling between a master and a slave peers.

The command supports the creation of an asynchronous mirroring coupling. Asynchronous mirroring is based on schedule-driven replication. The system also offers a predefined schedule object with a non-user-configurable interval of 20 seconds, named min_interval.

To create a mirroring coupling, an existing master peer must be specified together with a slave peer. Upon creation, the coupling is not active and the user needs to activate it explicitly in order to start the replication. This slave either already exists or is created by this command. Using an existing slave is allowed only if it is formatted. If the slave already exists, the command receives its name along with the remote system name. If it is created by this command, the input parameters specify the remote storage system name, the name of the slave that is created and the storage pool that will contain the newly created slave.

To add a second mirror (Multi-site) for an existing mirrored volume, use the part_of_multisite flag for the new mirror.

Mirroring is created in the standby state. The mirroring coupling must then be activated in order to start the initialization process, which copies the data from the master to the slave.

A storage system can have multiple mirroring definitions between pairs of peers on various remote systems. However, when the peers are consistency groups, all the volumes included in a specific consistency group must be mirrored between only one pair of storage systems. Therefore, when a volume peer on a storage system (for example: A) has a mirroring relation with a volume on a remote storage system (for example: B), any other volume in the same consistency group on storage system A can only be defined in a remote mirroring relation with a volume on storage system B. The same goes for volumes from storage system B to A. In addition, the mirrored consistency group has one sync job for all pertinent mirrored volumes within the consistency group.

The command fails if it finds conflicting mirroring snapshots (that were not removed during the deletion of a previous mirroring definition).

Initialization types:

- The **`online`** option (default)enables an over-the-wire initialization. In other words, it uses an inter-site link to replicate the master peer's initial state to the slave, starting once the mirror is first activated (mirror_activate). During initialization, the mirror status will be *Initialization*.
- If the **`offline`** option is selected, the initialization of the slave peer is not done by replicating the master's initial image, but rather by creating its offline replica. In other words, it restores to the slave a mirror image that is backed up on the master.

To create a remote mirroring coupling as part of a Multi-site relation, apply the part_of_multisite parameter.

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- VOLUME_SIZE_VERY_LARGE_ARE_YOU_SURE
  The volume size is very large. It may not be possible to mirror this volume to older versions of the storage system. Are you sure?

Return codes

- ASYNC_MIRROR_MISSING_RPO
  An asynchronous mirror definition must include the RPO.
- ASYNC_MIRROR_REMOTE_RPO_TOO_LONG
  The specified remote RPO is too long.
- ASYNC_MIRROR_REMOTE_RPO_TOO_SHORT
  The specified remote RPO is too short.
- ASYNC_MIRROR_RPO_TOO_SHORT
  The specified RPO is too short.
- ASYNC_MIRROR_RPO_TOO_LONG
  The specified RPO is too long.
- ASYNC_NOT_SUPPORTED_IN_TARGET
  The specified target does not support asynchronous mirroring.

- BAD_REMOTE_VOLUME_NAME
  The secondary volume name does not exist.
- BAD_REMOTE_VOLUME_SIZE
  The primary and secondary volumes contain a different number of blocks.
- CONS_GROUP_BAD_NAME
  The consistency group name does not exist.
- CONS_GROUP_HAS_MIRROR
  Mirroring is defined for this consistency group.
- CONS_GROUP_MIRRORING_NOT_SUPPORTED_IN_TARGET
  Consistency group mirroring is not supported by the target machine.
- INTERVAL_SHOULD_BE_SHORTER_THAN_RPO
  The schedule interval must be shorter than the RPO.
- ILLEGAL_INTERVAL
  The specified interval value is not supported.
- MAX_MIRRORS_REACHED
  The maximum number of mirrors is already reached.
- MAX_SYNC_MIRRORS_REACHED
  The maximum number of sync mirrors is already defined
- MAX_ASYNC_MIRRORS_REACHED
  The maximum number of async mirrors is already reached.
- NOT_ENOUGH_SPACE_ON_REMOTE_MACHINE
  Not enough free space to set the requested size of the secondary volume.
- NO_ASYNC_IN_THIN_PROVISIONED_POOL
  A thin-provisioned pool cannot contain volumes with asynchronous mirroring.
- VOLUME_BAD_NAME
  The volume name does not exist.
- VOLUME_IS_MASTER
  This local volume is already defined as a primary volume.
- VOLUME_IS_SLAVE
  The volume is defined as a secondary volume.
- REMOTE_VOLUME_EXISTS
  The secondary volume with the indicated name already exists. The name cannot be reused.
- REMOTE_MAX_VOLUMES_REACHED
  The maximum number of volumes on the remote machine is already reached.
- REMOTE_MAX_MIRRORS_REACHED
  The maximum number of mirrors is already defined on remote machine
- REMOTE_CONS_GROUP_NOT_EMPTY
  The remote consistency group contains volumes.
- VOLUME_BAD_PREFIX
  The volume name has a reserved prefix.
- REMOTE_POOL_DOES_NOT_EXIST
  The pool does not exist on the remote machine.
- REMOTE_POOL_NOT_SPECIFIED
  Prior to creating a secondary volume, a pool must be defined on the remote machine.
- REMOTE_TARGET_NOT_CONNECTED
  There is currently no connection from the target system.
- VOLUME_IS_SNAPSHOT
  The operation is not permitted on snapshots.
- REMOTE_VOLUME_IS_SNAPSHOT
  The secondary volume is a snapshot.
- TARGET_BAD_NAME
  The target name does not exist.
- TARGET_BAD_TYPE
  The target machine is not an XIV machine.
- TARGET_NO_ACCESS
  No access permissions to the secondary machine.

- TARGET_NOT_CONNECTED
  There is currently no connection to the target system.
- REMOTE_VOLUME_LOCKED
  The secondary volume is locked.
- TIMEOUT
  A remote operation was not completed in time.
- VOLUME_HAS_MIRRORING_SNAPSHOTS
  The volume has snapshots created by a previous mirroring process.
- SLAVE_VOLUME_NOT_FORMATTED
  The secondary volume is not formatted.
- TARGET_DOES_NOT_ACCEPT_XIV_COMMANDS
  The target system does not accept XIV management commands.
- SYNC_MIRROR_HAS_NO_RPO
  The synchronous mirror does not have an RPO.
- REMOTE_CONS_GROUP_IS_MIRRORED
  Mirroring is defined for this remote consistency group.
- REMOTE_SCHEDULE_DOES_NOT_EXIST
  The specified schedule does not exist on the remote machine.
- SCHEDULE_DOES_NOT_EXIST
  The specified schedule does not exist.
- REMOTE_CONS_GROUP_BAD_NAME
  The remote consistency group name does not exist.
- REMOTE_VOLUME_IS_MASTER
  A volume on the remote machine is already defined as primary.
- REMOTE_VOLUME_IS_SLAVE
  A volume on the remote machine is already defined as secondary.
- REMOTE_MAX_MIRROR_CAPACITY_REACHED
  The maximum capacity for mirrored volumes is already reached on the remote machine.
- MIRROR_RETRY_OPERATION
  There is an operation in progress on this mirror.

  Troubleshooting: Retry the command in a few seconds.
- MIRRORING_INCOMPATIBLE_TARGET_VERSION
  Mirroring is not supported between the system versions of the specified peers.
- MIRROR_TYPE_INCOMPATIBLE_WITH_TARGET
  A mirror of this type is not supported between the system versions of the specified peers.
- NO_OFFLINE_INIT_TYPE_WITH_SLAVE_CREATION
  A new volume will be created as secondary. Offline initialization is meaningless.
- ASYNC_WITH_OFFLINE_INIT_NOT_SUPPORTED_IN_TARGET
  The specified target does not support asynchronous mirroring with offline initialization.
- VOLUME_SIZE_ABOVE_LIMIT
  The specified volume size is above the limit.
- REMOTE_VOLUME_SIZE_ABOVE_LIMIT
  The specified volume size is above the limit of the remote machine.
- INVALID_SLICE_OFFSET
  Slice offset is illegal.
- VOLUME_IS_OLVM_PROXY
  The volume is in an IBM Hyper-Scale Mobility Proxy phase.
- REMOTE_VOLUME_IS_OLVM_PROXY
  The remote volume is in an IBM Hyper-Scale Mobility Proxy phase.
- ENCRYPTION_IN_PROGRESS
  The system is in the process of changing the encryption activation state.
- MIRROR_OF_SAME_TYPE_EXISTS_ON_VOLUME
  A mirror of this type is already defined on this volume.
- MIRROR_EXISTS_ON_TARGET
  The volume already has a mirror on this target.
- REMOTE_VOLUME_IS_MIRROR_MASTER

The volume is primary in a mirror relationship, and cannot be secondary!

- MULTISITE_MAX_NUM_OF_MIRRORS_REACHED
  Failed to create the relation, the maximum allowed number of relations is already exceeded.
- REMOTE_VOLUME_TWO_SYNC_MIRRORS_NOT_ALLOWED
  Two synchronous mirrors were detected on the remote volume. This is not allowed.
- REMOTE_VOLUME_MIRROR_LOOP_DETECTED
  A mirror loop was detected on the remote volume. This means that there is a mirror on the remote system, whose target is this system. Therefore, you cannot create a mirror with this target on this system.
- DOMAIN_MAX_MIRRORS_REACHED
  The domain exceeds the maximum allowed number of mirrors.
- REMOTE_DOMAIN_MAX_VOLUMES_REACHED
  The maximum number of volumes in the remote machine domain is already reached.
- REMOTE_DOMAIN_HAS_NO_ACCESS_TO_TARGET
  The secondary machine domain has no access to the target.
- REMOTE_DOMAIN_HAS_NO_ACCESS_TO_SCHEDULE
  The secondary machine domain has no access to the schedule.
- DOMAIN_HAS_NO_ACCESS_TO_TARGET
  The domain has no access to the target.
- REMOTE_DOMAIN_MAX_MIRRORS_REACHED
  The maximum number of mirrors is already reached in the remote machine domain.
- DOMAIN_MAX_VOLUMES_REACHED
  The domain exceeds the maximum allowed number of volumes.
- REMOTE_VOLUME_HAS_DATA_MIGRATION
  Data migration is already defined for the secondary volume.
- REMOTE_VOLUME_MASTER_ASYNC_MIRROR_DETECTED
  An asynchronous primary mirror was detected on the remote volume. The operation not allowed.
- MAX_MULTISITES_REACHED
  The number of Multi-site objects exceeded the limit.
- MULTISITE_INCOMPATIBLE_TARGET_VERSION
  Multi-site is not supported between the system versions of the specified peers.
- REMOTE_VOLUME_HAS_MIRRORING_SNAPSHOTS
  The remote volume has snapshots created by a previous mirroring process.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support
- SYSTEM_OUT_OF_PHYSICAL_SPACE
  The operation not allowed while the system is out of physical space.
- REMOTE_DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier of the remote system is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support
- REMOTE_SYSTEM_OUT_OF_PHYSICAL_SPACE
  The operation not allowed while the remote system is out of physical space.
- SLAVE_VOLUME_NOT_SAME_TYPE
  The primary and secondary volumes are not of the same type. Either one of them is compressed and the other is not, or they use different compression technologies.
- CONS_GROUP_IS_SECONDARY
  The consistency group is already defined as a secondary consistency group.
- CONS_GROUP_IS_PRIMARY
  The consistency group is already defined as a primary consistency group.
- RELATION_OF_SAME_TYPE_EXISTS_ON_CONS_GROUP
  A relation of this type is already defined on this consistency group.
- REMOTE_VOLUME_BAD_POOL
  The remote volume and remote consistency group belong to different storage pools.
- REMOTE_VOLUME_BELONGS_TO_CONS_GROUP
  The remote volume belongs to a consistency group.
- CONS_GROUP_MIRROR_ROLE_MISMATCH

All volumes in a mirrored consistency group must have the same mirroring role.

- CONS_GROUP_MIRROR_SCHEDULE_MISMATCH
  All volumes in a mirrored consistency group must have the same mirroring schedule.
- CONS_GROUP_RELATION_STATE_IS_NOT_INITIALIZING
  This operation is only allowed when the matching consistency group relation state is 'Initializing'.
- VOLUME_RELATION_IN_CG_CAN_NOT_BE_CREATED
  A second relation must first be added to the consistency group before it can be added to a member volume.
- CONS_GROUP_RELATION_PART_OF_MULTISITE_MISMATCH
  All volumes in a mirrored or HyperSwap consistency group must have the same part_of_multisite setting.
- CONS_GROUP_NOT_EMPTY
  This operation is only allowed on an empty consistency group.
- CONS_GROUP_MIRROR_TARGET_MISMATCH
  All volumes in a mirrored consistency group must have the same mirroring target.
- CONS_GROUP_RELATION_INIT_TYPE_MISMATCH
  All volumes in a mirrored or HyperSwap consistency group must have the same init type setting.
- REMOTE_MULTISITE_CONS_GROUP_MEMBER_VOL_HAS_TOO_MANY_RELATIONS
  A remote volume belonging to the consistency group has too many relations.
- MULTISITE_CANNOT_CONTAIN_SYNC_MIRROR
  A Multi-site relation cannot include a synchronous mirroring relation.
- REMOTE_VOLUME_HAS_A_PART_OF_MULTISITE_RELATION_BUT_NO_MULTISITE
  The target volume already has a 'part of Multi-site' relation. Creating a mirror relation on a volume with an existing 'part of Multi-site' relation is not allowed, unless the volume has a configured Multi-site relation.

  Troubleshooting: Define a Multi-site relation on this volume and retry the command.
- MIRROR_MULTISITE_MISMATCH
  The remote and local volumes are associated with different Multi-site relations.

Deactivating mirroring

Use the **mirror_deactivate** command to deactivate mirroring for a defined mirror coupling.

```
mirror_deactivate  < vol=<vol1[,vol2]...> | cg=cgName > [ target=TargetName ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	Master volume name or a list of master volumes.	N	N/A
cg	Object name	Master consistency group name or a list of master consistency groups.	N	N/A
target	Object name	Target mirror name. Mandatory if 2 mirrors are defined on the volume.	N	[none]

This command deactivates a coupling and switches it to the Inactive state. While in the Inactive state, only the master volume is updated, as opposed to the Active state, where the slave volume is updated together with the master volume.

The command cannot be issued on a slave.

If the mirroring is already inactive, this command has no effect and a success code is returned.

If more than one volume is specified, mirroring on all the volumes is deactivated. Furthermore, the deactivation of all the volumes is performed as an atomic operation, so that the slave volumes remain consistent with each other.

Deactivating a consistency group affects all of its volumes.

The command fails under the following conditions:

- The specified target does not exist.
- The specified target is non-mirrored.
- The specified target is a volume that belongs to a consistency group (in this case, the entire consistency group must be deactivated).
- Some of the specified targets are masters and some are slaves.
  - Each instance of the command can be applied to either master(s) or slave(s), but not to both.
- The target is a slave, yet the link is up.
- If multiple volumes are specified in the command and some are already part of an inactive mirror, the command will fail for all mirrors, including those that were active. The relevant return code is: SYNC_ALREADY_INACTIVE.

Access control

User Category	Permission
Storage administrator	Allowed

User Category	Permission
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.
- VOLUME_NO_MIRROR
  The local volume does not have remote mirroring definitions.
- CONS_GROUP_BAD_NAME
  The consistency group name does not exist.
- CONS_GROUP_NO_MIRROR
  The local consistency group does not have remote mirroring definitions.
- LOCAL_PEER_IS_NOT_MASTER
  The local peer is not primary.
- SYNC_ALREADY_INACTIVE
  Synchronization is already inactive.
- VOLUME_BELONGS_TO_MIRRORED_CONS_GROUP
  The volume mirror is part of a consistency group mirror.
- MIRROR_RETRY_OPERATION
  There is an operation in progress on this mirror.

  Troubleshooting: Retry the command in a few seconds.
- TARGET_BAD_NAME
  The target name does not exist.
- VOLUME_HAS_MULTIPLE_MIRRORS
  The volume has multiple mirrors. The operation is not allowed or a target must be specified.
- CONS_GROUP_HAS_MULTIPLE_MIRRORS
  The consistency group has multiple mirrors. The operation is not allowed, or a target must be specified.
- VOLUME_TARGET_MISMATCH
  The volume and target do not match.
- CONS_GROUP_BAD_TARGET
  The target name does not match the consistency group.
- REMOTE_MIRROR_IS_STANDBY
  The remote mirror is marked as Standby.
- MIRROR_IS_STANDBY
  The mirror is marked as Standby.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support

Deleting a remote mirroring definition

Use the **mirror_delete** command to delete a remote mirroring coupling definition.

```
mirror_delete < vol=VolName | cg=cgName > [ target=TargetName ] [ force_on_slave=<Yes|No> ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	Local master volume name.	N	N/A
cg	Object name	Local master consistency group name.	N	N/A
target	Object name	Target mirror name. Mandatory if 2 mirrors are defined on the volume.	N	[none]

Name	Type	Description	Mandatory	Default
force_on_slave	Boolean	Forces the deletion of the remote mirroring coupling definition even of a slave. Deleting a remote mirroring definition can be forced on the slave peer only when it is in the initialization phase.	N	no

When a coupling is initially created or after it is deactivated, it is in *standby* mode. Only a standby coupling can be deleted. The command can only be issued on the master.

After the remote mirroring is deleted, both peers are configured as *none*, meaning that they are no longer configured as either master or slave.

Only the remote mirroring coupling definition is deleted. Neither the volumes themselves, nor their snapshots are deleted.

The local object specified in the vol parameter, must be a master.

To delete a remote mirroring coupling, the communication must be established. If there is no communication, mirroring is only deleted on the master, and a configuration error appears on the slave once the communication resumes.

Command outcome:

- An event is generated
- Overall coupling statistics are captured
- The outstanding pertinent sync jobs are deleted
- The process completion is recorded in the log

Deleting the mirroring definition when the link is down:

- When the link is down, this command only deletes the mirroring definition on the master.
- To delete the mirroring definition from the slave:
  - Run the mirror_change_role command to turn the slave into the master
  - Run mirror_delete

The force_on_slave parameter:

- The parameter force_on_slave can be issued only if mirroring is in the initialization phase. In any other mode, the role can be changed to master and the peer mirror can be deleted.

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_DELETE_CG_MIRRORING
  Are you sure you want to delete the mirroring relations of the consistency group and of all volumes in the consistency group?

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.
- VOLUME_NO_MIRROR
  The local volume does not have remote mirroring definitions.
- CONS_GROUP_BAD_NAME
  The consistency group name does not exist.
- CONS_GROUP_NO_MIRROR
  The local consistency group does not have remote mirroring definitions.
- LOCAL_PEER_IS_NOT_MASTER
  The local peer is not primary.
- MIRROR_IS_ACTIVE
  Remote mirroring is currently active.
- FORCE_DELETE_NOT_ALLOWED_ON_MASTER
  Deletion needs to be forced on secondary mirrors only.
- VOLUME_BELONGS_TO_MIRRORED_CONS_GROUP
  The volume mirror is part of a consistency group mirror.
- MIRROR_RETRY_OPERATION

There is an operation in progress on this mirror.

Troubleshooting: Retry the command in a few seconds.

- MIRROR_IS_NOT_INITIALIZING
  The operation is permitted only during the Initialization phase.
- TARGET_BAD_NAME
  The target name does not exist.
- VOLUME_HAS_MULTIPLE_MIRRORS
  The volume has multiple mirrors. The operation is not allowed or a target must be specified.
- CONS_GROUP_HAS_MULTIPLE_MIRRORS
  The consistency group has multiple mirrors. The operation is not allowed, or a target must be specified.
- VOLUME_TARGET_MISMATCH
  The volume and target do not match.
- MIRROR_ASSOCIATED_WITH_MULTISITE
  This mirror is associated with a defined Multi-site, the operation is not allowed.
- CONS_GROUP_BAD_TARGET
  The target name does not match the consistency group.
- REMOTE_MIRROR_IS_STANDBY
  The remote mirror is marked as Standby.
- MIRROR_IS_STANDBY
  The mirror is marked as Standby.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support
- MIRROR_CONS_GROUP_MEMBERSHIP_MISMATCH
  The mirrored consistency group contains different volumes on the primary and secondary machines. This problem occurs whenever the cg_add_vol or cg_remove_vol commands were previously issued, and the primary machine did not receive an acknowledgment from the secondary machine until the command timed out, or due to any other unexpected failure.

Viewing the mirroring status

Use the **mirror_list** command to list the status and configuration of mirroring couplings.

```
mirror_list [ < [ vol=VolName ] [ target=TargetName ] > | cg=cgName | < [ scope=<cg|volume> ]
 [ sync_type=<sync_best_effort|async_interval> ] > ] [ domain=DomainName ]
```

Parameters

Name	Type	Description	Mandatory	Default
sync_type	Enumeration	List type. The avaiable options are: **SYNC_BEST_EFFORT**, **ASYNC_INTERVAL**, or All (if no value is specified)	N	All (if no value is specified)
scope	Enumeration	List type: all mirrors, all volumes, all CGs	N	All (if no value is specified)
vol	Object name	Local volume name.	N	[none]
cg	Object name	Local consistency group name.	N	[none]
target	Object name	Remote target name.	N	[none]
domain	Object name	The domain name.	N	All Domains

This command shows current configuration and status for the remote mirroring of volumes or consistency groups. Size/part/time to synchronize are unknown if this is the slave and connection is broken.

The following default parameters are shown:

- Name
- Mirror Type (**sync_best_effort**, **async_interval**)
- Mirror Object: **CG** or **Volume**
- Role: **Master** or **Slave**
- Remote System: target name
- Remote Peer: volume name
- Active: **Yes** or **No**
- Status: **Initializing**, **Synchronized**, **Unsynchronized**, **Consistent**, **Inconsistent**, **RPO OK**, **RPO Lagging**, or **Change Tracking**
- Link Up: **Yes** or **No**

The following optional parameters can be listed by explicitly specifying the proper columns:

- Designation: **Primary** or **Secondary**

- Estimated Sync Time: Estimated time to synchronization in seconds
- Size To Synchronize: Size to synchronize in MiB
- Operational: **Yes** or **No**
- Sync Progress (in %)
- Mirror Error: Explains why mirroring is deactivated: **No_Error**, **Configuration_Error**, **Secondary_Pool_Exhausted**, **Master_Pool_Exhausted**, or **No_Thin_Provisioning_Resources**
- Schedule Name
- Last replicated Snapshot Time: the value is presented in yyyy-mm-dd hh:mm:ss format
- Specified RPO: the value is presented in h:mm:ss format

The following deactivation reasons can be read from the output list (available only in XML output format):

- **No_Error**
- **Configuration_Error**
- **Secondary_Pool_Exhausted**
- **Master_Pool_Exhausted**
- **Remote_And_Local_Volume_Size_Mismatch**
- **Cons_Group_Membership_Mismatch**
- **Possible_Remote_And_Local_Volume_Size_Mismatch**
- **Possible_Cons_Group_Membership_Mismatch**
- **No_Thin_Provisioning_Resources**
- **Peer_Status_Mismatch**
- **Temporarily_Deactivated_For_Upgrade**
- **Mirror_is_Standby**
- **Out_Of_Physical_Space**
- **Data_Reduction_Offline**
- **Secondary_Out_Of_Physical_Space**
- **Secondary_Data_Reduction_Offline**
- **HA_Slave_Failover**
- **Data_Service_failure**
- **Remote_Volume_Down**
- **SCSI_Reservation_Error**

Field ID	Field output	Description	Default position
local_peer_name	Name	N/A	1
mirror_object	Mirror Object	N/A	3
designation	Designation	N/A	N/A
current_role	Role	N/A	4
target_name	Remote System	N/A	5
remote_peer_name	Remote Peer	N/A	6
active	Active	N/A	7
sync_state	Status	N/A	9
connected	Link Up	N/A	10
estimated_sync_time	Est. remaining time (sec)	N/A	N/A
size_to_synchronize	Size To Sync (MiB)	N/A	N/A
operational	Operational	N/A	N/A
sync_progress	Sync Progress (%)	N/A	N/A
mirror_error	Mirror Error	No Error, Secondary pool exhausted, Configuration error or No thin provisioning resources	N/A
sync_type	Mirror Type	N/A	2
schedule_name	Schedule Name	N/A	N/A
last_replicated_snapshot_time	Last Replicated	N/A	N/A
last_replicated_snapshot_exists	Has Last Replicated Snapshot	N/A	N/A
specified_rpo	RPO	N/A	N/A
remote_rpo	Remote RPO	N/A	N/A
crash_consistent	Crash Consistency	N/A	N/A
validate	Validation	N/A	N/A
is_standby	Standby	N/A	8
arch	Remote Arch	N/A	N/A

Output:

```
<command id="_home_markdown_jenkins_workspace_Transform_in_STJKN5_12.3.2_CLI_latest_r_command__mirroring__mirror_list_0">
<administrator>
    <command>
        <changes_session_id value="1288716489394201:1:1288903896317961:1"/>
        <code value="SUCCESS"/>
        <last_change_index value="32289"/>
        <status value="0"/>
        <status_str value="Command completed successfully"/>
        <return>
            <mirror
id="_home_markdown_jenkins_workspace_Transform_in_STJKN5_12.3.2_CLI_latest_r_command__mirroring__mirror_list_100777">
```

```
                <id value="100777"/>
                <creator value=""/>
                <creator_category value="none"/>
                <local_peer_id value="100776"/>
                <local_peer_name value="SYNC_vol_5"/>
                <schedule_name value=""/>
                <designation value="Secondary"/>
                <current_role value="Slave"/>
                <remote_mirror_id value="100872"/>
                <remote_peer_name value="SYNC_vol_4"/>
                <target_id value="100707"/>
                <target_name value="SYNC_target_2"/>
                <sync_type value="sync_best_effort"/>
                <sync_state value="Consistent"/>
                <active value="yes"/>
                <connected value="yes"/>
                <operational value="yes"/>
                <sync_progress value="100"/>
                <size_to_synchronize value="-1"/>
                <estimated_sync_time value="0"/>
                <mirror_error value="No_Error"/>
                <mirror_object value="Volume"/>
                <specified_rpo value=""/>
                <remote_rpo value=""/>
                <last_replicated_snapshot_time value=""/>
                <init_type value="online"/>
            </mirror>
        </return>
    </command>
</administrator>
<aserver status="DELIVERY_SUCCESSFUL"/>
</command>
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Obtaining statistics on past sync jobs

Use the **mirror_statistics_get** command to present statistics that are automatically gathered by the system on past sync jobs per specified mirrored volume or consistency job.

```
mirror_statistics_get <vol=VolName | cg=cgName> [ target=TargetName ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	Local volume name.	N	N/A
cg	Object name	Local consistency group name.	N	N/A
target	Object name	Target mirror name. Mandatory if 2 mirrors are defined on the volume.	N	[none]

The command output includes:

- Date and time created
- Date and time started to run
- Date and time finished
- Job size (MiB)

Either a volume or consistency group must be specified.

Field ID	Field output	Default position
created_at	Created	1
started_at	Started	2
finished_at	Finished	3
job_size	Job Size (MiB)	4
duration	Job Duration (Sec)	5
avg_sync_rate	Average Sync Rate (MB/sec)	6

Example:

```
mirror_statistics_get vol=VolName
```

Output:

```
<job
id="_home_markdown_jenkins_workspace_Transform_in_STJKN5_12.3.2_CLI_latest_r_command__mirroring__mirror_statistics_get_14
3">
    <avg_sync_rate value="22.3333"/>
    <created_at value="2011-03-22 11:19:30"/>
    <duration value="6"/>
    <finished_at value="2011-03-22 11:19:36"/>
    <job_size value="134"/>
    <started_at value="2011-03-22 11:19:30"/>
</job>
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.
- CONS_GROUP_NO_MIRROR
  The local consistency group does not have remote mirroring definitions.
- MIRROR_HAS_NO_STATISTICS
  Job statistics were not collected for this mirror.
- LOCAL_IS_SLAVE
  The local mirror peer is not primary.
- VOLUME_BELONGS_TO_MIRRORED_CONS_GROUP
  The volume mirror is part of a consistency group mirror.
- VOLUME_NO_MIRROR
  The local volume does not have remote mirroring definitions.
- CONS_GROUP_BAD_NAME
  The consistency group name does not exist.
- SYNC_MIRROR_HAS_NO_STATISTICS
  Job statistics do not exist for the synchronous mirror.
- TARGET_BAD_NAME
  The target name does not exist.
- VOLUME_HAS_MULTIPLE_MIRRORS
  The volume has multiple mirrors. The operation is not allowed or a target must be specified.
- CONS_GROUP_HAS_MULTIPLE_MIRRORS
  The consistency group has multiple mirrors. The operation is not allowed, or a target must be specified.
- VOLUME_TARGET_MISMATCH
  The volume and target do not match.
- CONS_GROUP_BAD_TARGET
  The target name does not match the consistency group.
- MIRROR_IS_STANDBY
  The mirror is marked as Standby.

Switching roles between master and slave

Use the **mirror_switch_roles** command to switch roles between master and slave volumes.

```
mirror_switch_roles <vol=VolName | cg=cgName> [ target=TargetName ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	Local volume name.	N	N/A
cg	Object name	Local consistency group name.	N	N/A
target	Object name	N/A	N	[none]

The command can only be issued if coupling is operational and only on the master. For synchronous mirroring it can only be issued when the coupling is synchronized; for asynchronous mirroring it can only be issued if there are no outstanding sync jobs and the volume and its last replicated snapshot are identical.

Following the execution of the command:

- The volume that was previously the master becomes the slave
- The volume that was previously the slave becomes the master

Before this command switches roles, the system stops accepting new writes to the local volume. With synchronous mirrors the system performs all pending writes, and only after all pending writes have been committed, the roles are switched.

After the command is executed, the mirror remains active.

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_SWITCH_ROLES
  Are you sure you want to switch the roles in this relation?

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.
- VOLUME_NO_MIRROR
  The local volume does not have remote mirroring definitions.
- CONS_GROUP_BAD_NAME
  The consistency group name does not exist.
- CONS_GROUP_NO_MIRROR
  The local consistency group does not have remote mirroring definitions.
- LOCAL_PEER_IS_NOT_MASTER
  The local peer is not primary.
- MIRROR_IS_NON_OPERATIONAL
  The mirror is non-operational.
- MIRROR_IS_NOT_SYNCHRONIZED
  The mirror is not synchronized.
- VOLUME_HAS_DATA_MIGRATION
  Data Migration is defined for this volume.
- REMOTE_TARGET_NOT_CONNECTED
  There is currently no connection from the target system.
- VOLUME_BELONGS_TO_MIRRORED_CONS_GROUP
  The volume mirror is part of a consistency group mirror.
- MIRROR_HAS_SYNC_JOB
  The operation is not permitted on a mirror with active sync jobs.
- MIRROR_RETRY_OPERATION
  There is an operation in progress on this mirror.

  Troubleshooting: Retry the command in a few seconds.
- MIRROR_MASTER_DIFFERS_FROM_SLAVE
  The mirror's primary volume was written to after the last replicated snapshot was taken.
- REMOTE_MIRROR_IS_NOT_ACTIVE
  Remote mirroring is currently inactive.

- TARGET_BAD_NAME
  The target name does not exist.
- VOLUME_HAS_MULTIPLE_RELATIONS
  The volume has multiple relations. The operation is not allowed or a target must be specified.
- CONS_GROUP_HAS_MULTIPLE_RELATIONS
  The consistency group has multiple relations. The operation is not allowed or a target must be specified.
- VOLUME_TARGET_MISMATCH
  The volume and target do not match.
- CONS_GROUP_BAD_TARGET
  The target name does not match the consistency group.
- REMOTE_MIRROR_IS_STANDBY
  The remote mirror is marked as Standby.
- MIRROR_IS_STANDBY
  The mirror is marked as Standby.
- COMMAND_NOT_SUPPORTED_FOR_OLVM_VOLUMES
  This command is not supported for IBM Hyper-Scale Mobility volumes.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support
- SYSTEM_OUT_OF_PHYSICAL_SPACE
  The operation not allowed while the system is out of physical space.

Retrieving RPO thresholds

Use the **rpo_thresholds_get** command to list system RPO-related thresholds, that, once crossed, trigger the creation of a corresponding event.

```
rpo_thresholds_get
```

Field ID	Field output	Default position
increase_percentage	Increase Percentage	1
increase_absolute	Increase Absolute	2

Example:

```
rpo_thresholds_get
```

Output:

```
Increase Percentage   Increase Absolute
--------------------- -------------------
100                   3600
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Setting an RPO threshold

Use the **rpo_thresholds_set** command to set system RPO-related thresholds, that, once crossed, trigger the creation of a corresponding event.

```
rpo_thresholds_set  [ increase_percentage=percentage ] [ increase_absolute=absolute ]
```

Parameters

Name	Type	Description	Mandatory	Default

Name	Type	Description	Mandatory	Default
increase_percentage	Integer	The threshold for RPO increase (in per cent), beyond which an event should be created.	N	none
increase_absolute	Integer	The threshold for RPO increase, beyond which an event should be created.	N	none

Example:

```
rpo_thresholds_set increase_percentage=percentage
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- INVALID_RPO_THRESHOLD_PERCENTAGE
  The values should be in the [1,10000] range.
- INVALID_RPO_THRESHOLD_ABSOLUTE
  The values should be in the [1,1000000] range.

Changing the interval of a schedule

Use the **schedule_change** command to change the interval of a schedule.

```
schedule_change schedule=Schedule interval=IntervalSize [ domain=DomainList ]
```

Parameters

Name	Type	Description	Mandatory	Default
schedule	Object name	The name of the schedule.	Y	N/A
interval	N/A	The interval for asynchronous mirroring. Format: hh:mm [:ss].	Y	N/A
domain	N/A	The schedule will be attached to the specified domains. To specify several domains, separate them with a comma. To specify all existing domains, use "*".	N	none

This command updates the schedule definition. Such definition can be referenced to when specifying asynchronous mirroring couplings.

Limitation:

- Only the following values are allowed in a schedule: 00:00:20, 00:00:30, 00:00:40, 00:00:50, 00:01, 00:01:10, 00:01:20, 00:01:30, 00:01:40, 00:01:50, 00:02, 00:05, 00:10.
- A predefined schedule cannot be changed.

Outcome:

- If the update command is issued on a schedule that is not referenced by any object, a confirmation message is displayed.
- If the update command is issued on a schedule that is referenced to by an object (for example, mirroring couplings), a warning message is displayed.
- Sync jobs that are running will not be affected.
- Future sync jobs are scheduled based on the new schedule settings.

Example:

```
schedule_change interval=00:01 schedule=1min domain=* -y
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed

User Category	Permission
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_UPDATE_THE_SCHEDULE
  Are you sure you want to update this schedule? This change will affect all the mirrors using that schedule.

Return codes

- SCHEDULE_DOES_NOT_EXIST
  The specified schedule does not exist.
- BAD_SCHEDULE_TIME_FORMAT
  Time format for a schedule is HH:MM[:SS].
- ILLEGAL_INTERVAL
  The specified interval value is not supported.
- SCHEDULE_CAN_NOT_BE_UPDATED
  The specified schedule cannot be updated.
- INTERVAL_SCHEDULE_REQUIRES_ONLY_ONE_INTERVAL
  Only one interval can be defined in an interval schedule.
- SCHEDULE_EXCLUDE_TIMES_NOT_REQUIRED
  An exclusion period can be defined only if exclude_time is set.
- ZERO_LENGTH_EXCLUSION_PERIOD
  An exclusion period's start time must be different from its end time.
- DOMAIN_SCHEDULE_IN_USE
  The schedule is in use, and therefore cannot be moved to another domain.
- DOMAIN_DOESNT_EXIST
  The domain does not exist.
- ILLEGAL_TIME_SLOT_SPACE
  The defined value must be larger than the minimum time slot, but smaller than the interval.

Creating a schedule object

Use the **schedule_create** command to define a schedule for replication.

```
schedule_create schedule=Schedule [ interval=IntervalSize ] [ type=<manual|interval|max|time> ]
 [ domain=DomainList ]
```

Parameters

Name	Type	Description	Mandatory	Default
schedule	Object name	The name of the schedule	Y	N/A
interval	N/A	The interval for asynchronous mirroring. Format: hh:mm [:ss].	N	00:10[:00]
type	Enumeration	The schedule type for asynchronous mirroring. Can be Manual or Interval.	N	interval
domain	N/A	The schedule will be attached to the specified domains. To specify several domains, separate them with a comma. To specify all existing domains, use "*".	N	none

This command creates a schedule definition. Schedules can be referenced to when specifying asynchronous mirroring couplings.

Limitations:

- Only the following values are allowed in a schedule: 00:00:20, 00:00:30, 00:00:40, 00:00:50, 00:01, 00:01:10, 00:01:20, 00:01:30, 00:01:40, 00:01:50, 00:02, 00:05, 00:10.
- The system features a predefined schedule object with a non-user-configurable interval of 20 seconds, named min_interval.

The type parameter:

Prior to the introduction of this parameter, each asynchronous mirror could be configured with an automatic schedule, whose interval specified how often a replication point and the corresponding replication process (sync job) should be automatically created. It was also possible to instruct the system to create a manual replication point and a corresponding sync job for a mirror using the dedicated CLI command mirror_create_snapshot. Finally, a single predefined schedule named *Never* with no interval settings was provided for mirrors that only required manual sync job creation.

The type parameter enables you to define multiple custom, user-configurable manual schedules. The creation of consistent, identical replication points for all mirrors set with such schedule, as well as corresponding sync jobs can be triggered using the dedicated CLI command schedule_create_tick, that specifies the

schedule name as an argument. This facilitates external/scripted replication control for mirrors sharing the same schedule, without requiring them to be interval-based.

When type=interval, synchronization jobs for a mirror associated with the schedule will be triggered automatically, based on the specified interval.

When type=manual, synchronization jobs for a mirror associated with the schedule can be triggered by the command schedule_create_tick.

Once set, the schedule type cannot be changed.

Example:

```
schedule_create interval=00:01 schedule=1min domain=*
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- SCHEDULE_EXISTS
  The schedule name exists.
- BAD_SCHEDULE_TIME_FORMAT
  Time format for a schedule is HH:MM[:SS].
- MAX_SYNC_SCHEDULES_REACHED
  The maximum number of schedule objects has been reached.
- ILLEGAL_INTERVAL
  The specified interval value is not supported.
- INTERVAL_SCHEDULE_REQUIRES_ONLY_ONE_INTERVAL
  Only one interval can be defined in an interval schedule.
- ZERO_LENGTH_EXCLUSION_PERIOD
  An exclusion period's start time must be different from its end time.
- SCHEDULE_EXCLUDE_TIMES_NOT_REQUIRED
  An exclusion period can be defined only if exclude_time is set.
- ONLY_INTERVAL_SCHEDULE_MAY_HAVE_EXCLUSIONS
  An exclusion period may defined only for an interval schedule.
- DOMAIN_DOESNT_EXIST
  The domain does not exist.
- ILLEGAL_TIME_SLOT_SPACE
  The defined value must be larger than the minimum time slot, but smaller than the interval.

Triggering a schedule

Use the **schedule_create_tick** command to trigger a schedule-equivalent event for the couplings with the specified schedule.

```
schedule_create_tick schedule=Schedule
```

Parameters

Name	Type	Description	Mandatory
schedule	Object name	The name of an asynchronously mirrored schedule.	Y

This command triggers a schedule-equivalent, interval-arrived event for couplings with the specified schedule.

- The command triggers a new sync job for asynchronous mirror definitions that are configured with the manual schedule specified by the command. The command triggers a simultaneous event for all mirrors with the specified schedule (and only whenever the schedule is of a non-interval type) which is equivalent to the 'new-interval-arrived' event triggered automatically by the system for a mirror (with a schedule of type interval).

- The command is different from mirror_create_snapshot whereas it is applied to mirrors that do not have an interval-based schedule. Thus, even though an event is triggered immediately (as with mirror_create snapshot), no sync job is created for a pertinent mirror with the specified schedule (in case such a mirror has an outstanding sync job, as one might expect for mirrors with an interval-based schedule, if a new interval arrives during an outstanding job).
- The event is triggered for all pertinent couplings at the same time.
- A warning is displayed, requiring a user confirmation.

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- SCHEDULE_DOES_NOT_EXIST
  The specified schedule does not exist.
- SCHEDULE_IS_NOT_MANUAL
  The specified schedule does not allow an external trigger.

Deleting a schedule object

Use the **schedule_delete** command to delete a schedule for replication.

```
schedule_delete schedule=Schedule
```

Parameters

Name	Type	Description	Mandatory
schedule	Object name	The name of the schedule to be deleted.	Y

This command deletes a schedule definition.

The command can be issued successfully only if the schedule specified is not referenced by a mirror coupling, or if it is not a pre-defined schedule (min_interval).

Outcome:

- The command will delete the specified schedule.

Example:

```
schedule_delete schedule=hourly
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- SCHEDULE_IS_ASSIGNED
  The specified schedule is currently assigned to a mirror.
- SCHEDULE_CAN_NOT_BE_DELETED
  The specified schedule cannot be deleted.
- SCHEDULE_DOES_NOT_EXIST
  The specified schedule does not exist.

Listing a schedule object

Use the **schedule_list** command to list the schedule properties for the specified coupling.

```
schedule_list [ schedule=Schedule ] [ domain=DomainName ]
```

Parameters

Name	Type	Description	Mandatory	Default
schedule	Object name	The name of the schedule.	N	All
domain	Object name	The domain name.	N	All Domains

The following default parameters are listed:

- Name
- Interval

The following optional parameters can be listed:

- Predefined (is the schedule a predefined object)
- Last Tick (last timestamp the schedule was fired)

Field ID	Field output	Default position
name	Name	1
interval	Interval	2
predefined	Predefined	N/A

Example:

```
schedule_list
```

Output:

```
Name            Interval
never
min_interval    00:00:20
ASYNC_None_3    00:02:00
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Renaming a schedule

Use the **schedule_rename** command to rename a schedule object.

```
schedule_rename schedule=Schedule new_name=Name
```

Parameters

Name	Type	Description	Mandatory
schedule	Object name	The current name of the schedule.	Y
new_name	Object name	The new name for the schedule.	Y

It is not possible to rename a predefined schedule.

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- SCHEDULE_DOES_NOT_EXIST
  The specified schedule does not exist.
- SCHEDULE_NAME_EXISTS
  The new schedule name already exists.
- SCHEDULE_CAN_NOT_BE_UPDATED
  The specified schedule cannot be updated.

Viewing sync job status

Use the **sync_job_list** command to list the statuses of queued and running sync jobs for asynchronous couplings

```
sync_job_list [ vol=VolName | cg=cgName ] [ domain=DomainName ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	Local volume name.	N	[none]
cg	Object name	Local consistency group name.	N	[none]
domain	Object name	The domain name.	N	All Domains

The following parameters are displayed:

- Mirroring coupling (volume/consistency group)
- Job state (provides state of the sync job transmission process, running from the primary to secondary peer):
  - new: sync job is initiating
  - ready: sync job is ready to transfer from primary to secondary peer
  - active: sync job is transferring newly written date to secondary peer from primary peer
  - done: sync job is complete - appears only for volumes whose mirror sync jobs have completed but under a consistency group mirror sync job that is still active
- Job type: scheduled, ad hoc
- Schedule (name of the referenced schedule object)
- Interval length (if applicable)
- Job size
- Job progress
- Date created
- Time created
- Date started to run
- Time started to run

Field ID	Field output	Default position
job_object	Job Object	1
mirror_peer	Local Peer	2
source_snap	Source	3
target_snap	Target	4
job_state	State	5
part_of_cg_job	Part of CG	6
job_type	Job Type	7
created_at	Created	N/A
started_at	Started	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

HyperSwap commands

This chapter describes the command line interface (CLI) for HyperSwap.

HyperSwap® delivers highly-available, non-disruptive storage service, through partial or complete system failures and disasters, in the same data center and between metro-distant data centers.

HyperSwap high availability is based on active-active pairing of storage systems per volume or per consistency group. Each volume or consistency group pair uses synchronous replication to keep both systems updated at all times.

When certain conditions apply, an automatic and completely transparent failover is performed, so that the applications experience no downtime. As soon as the actual failure is recovered, the pair is automatically resynchronized.

As in other high availability solutions, HyperSwap requires a quorum witness component, to avoid split-brain situations. HyperSwap Quorum Witness is constantly monitoring the status of the related storage systems, and, if necessary, acts as a tiebreaker for conflict resolution.

The HyperSwap solution relies on Asymmetrical Logical Unit Access (ALUA) support to inform the host about the optimized paths to the storage system, and minimize I/O latency.

FlashSystem A9000 and FlashSystem A9000R HyperSwap capability does not require additional special hardware or software, and does not require any additional licensing.
Important: The input and output syntax of CLI commands uses the legacy terminology of "Master", "SMaster", and "Slave" volumes, which in any documentation except the CLI reference, are referred to as "Primary", "Secondary", and "Tertiary". This inconsistency is a necessary compromise, required to avoid changes to older CLI commands that are in customer use, and also to keep the CLI terminology consistent across the board. The new terminology helps emphasize the commonality between the more recent functions of Multi-site HA/DR, high availability (HyperSwap), and the disaster recovery (Synchronous and Asynchronous mirroring) ones. It is used outside the CLI reference, where broader concepts can be explained.

- **Creating a HyperSwap relation**
  Use the **ha_create** command to create a HyperSwap relation.
- **Viewing the status of HyperSwap volumes and consistency groups**
  Use the **ha_list** command to display the status of HyperSwap volumes and consistency groups.
- **Activating a HyperSwap relation**
  Use the **ha_activate** command to activate a HyperSwap relation of volumes or consistency groups.
- **Deactivating a HyperSwap relation**
  Use the **ha_deactivate** command to deactivate HyperSwap volumes or consistency groups.
- **Deleting a HyperSwap relation**
  Use the **ha_delete** command to delete a HyperSwap relation.
- **Switching roles between Master and Slave volumes**
  Use the **ha_switch_roles** command to switch roles between Master and Slave volumes.
- **Changing a peer role in a HyperSwap volume**
  Use the **ha_change_role** command to change the role of a local HyperSwap relation peer from Master to Slave or from Slave to Master.
- **Restoring the availability of a Master volume**
  Use the **ha_restore_availability** command to restore the availability of a Master volume, that became unavailable due to a failure.
- **Creating a HyperSwap volume snapshot (ad hoc sync job)**
  Use the **ha_create_snapshot** command to simultaneously create snapshots on both peers of a HyperSwap relation.
- **Changing the designation of HyperSwap relation peers**
  Use the **ha_change_designation** command to change the designation of HyperSwap relation peers from Primary to Secondary and vice versa.
- **Enabling high availability of a HyperSwap relation**
  Use the **ha_high_availability_enable** command to enable high availability (automatic failover) of a HyperSwap relation.
- **Disabling high availability of a HyperSwap relation**
  Use the **ha_high_availability_disable** command to disable the high availability (automatic failover) of a HyperSwap relation.
- **Converting a HyperSwap relation into a sync mirror**
  Use the **ha_convert_into_mirror** command to change a HyperSwap relation into a sync mirror.
- **Converting a sync mirror into a HyperSwap relation**
  Use the **mirror_convert_into_ha** command to change a sync mirror into a HyperSwap relation.
- **Creating a new Quorum Witness**
  Use the **quorum_witness_define** command to create a new Quorum Witness definition in the system and connect the system to the Quorum Witness.
- **Listing Quorum Witnesses**
  Use the **quorum_witness_list** command to list all Quorum Witnesses defined in the system, or only the specified one.
- **Updating a Quorum Witness definition**
  Use the **quorum_witness_update** command to update a Quorum Witness definition.
- **Renaming a Quorum Witness**
  Use the **quorum_witness_rename** command to rename a Quorum Witness.
- **Deleting a Quorum Witness**
  Use the **quorum_witness_delete** command to delete a Quorum Witness.
- **Activating a Quorum Witness**
  Use the **quorum_witness_activate** command to activate a Quorum Witness.
- **Deactivating a Quorum Witness**
  Use the **quorum_witness_deactivate** command to deactivate a Quorum Witness.
- **Listing the Quorum Witness connection status**
  Use the **quorum_witness_connections_list** command to list the status of the connections with the Quorum Witness.
- **Getting Quorum Witness information**
  Use the **quorum_witness_info_get** command to get Quorum Witness information from a Quorum Witness.

Creating a HyperSwap relation

Use the **ha_create** command to create a HyperSwap relation.

```
ha_create < vol=VolName [ create_slave=<yes|no> [ remote_pool=RemotePoolName ] ]
 [ init_type=<online|offline> ]
```

```
> | <cg=cgName slave_cg=SlaveCgName> target=TargetName  [ part_of_multisite=<yes|no> ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	Local volume to be replicated (the Master).	N	N/A
create_slave	Boolean	Determines whether to create a new Slave volume or to use an existing one. If an existing remote volume is used, its name must match the local volume name.	N	no
remote_pool	Object name	The storage pool on the remote system. Relevant only if creating a Slave volume.	N	N/A
cg	Object name	The local consistency group to be mirrored.	N	N/A
slave_cg	Object name	The name of the Slave consistency group on the remote storage system.	N	N/A
target	Object name	The remote target to contain the Slave volume.	Y	N/A
init_type	Enumeration	The initialization method of the Slave volume.	N	[none]
part_of_multisite	Boolean	Marks the HA as part of Multi-site.	N	no

This command defines a new HyperSwap relation between Master and Slave peers. When you define a HyperSwap relation, the following rules are enforced:

- The Slave volume is not mapped
- The Master and Slave peers must be configured with the same QW
- The Master and Slave connectivity to the QW must be healthy
- If the HyperSwap relation includes an existing Slave volume, the Slave volume's name must match the Master volume's name

As part of the operation, Master volume metadata is copied to the Slave volume, and the Slave volume identity changes.

A HyperSwap relation is created in the Standby state. It must then be activated in order to start the initialization process, which copies data from the Master to the Slave.

The following initialization methods are available:

- The online option (default) enables an over-the-wire initialization. In other words, it uses an inter-site link to replicate the Master's initial state to the Slave, starting once the mirror is first activated (ha_activate). During initialization, the HyperSwap relation status will be Initializing.
- If the offline option is selected, the initialization of the Slave peer is not done by replicating the Master's initial image, but rather by creating its offline replica. In other words, it restores to the Slave a mirror image that is backed up on the Master. Once the relation is activated, the contents of the volumes are compared, and only modified data is synchronized over the wire. This process is usually much faster than online initialization. During initialization, the HyperSwap relation status will be Initializing.

A storage system can have multiple HyperSwap relations between pairs of peers on various remote systems. However, when the peers have consistency group HyperSwap relations, all the volumes included in a specific consistency group HyperSwap relation can only be replicated between one pair of storage systems. Therefore, when a volume peer on a storage system (for example: A) has a HyperSwap relation with a volume on a remote storage system (for example: B), any other volume in the same consistency group on storage system A can only be defined in a HyperSwap relation with a volume on storage system B. The same is true for volumes from storage system B to A. In addition, the HyperSwap consistency group has one sync job for all pertinent HyperSwap volumes within the consistency group. To create a HyperSwap relation as part of a Multi-site relation, apply the part_of_multisite parameter.

Example:

```
ha_create vol=regular_volume target=123 create_slave=yes remote_pool=333
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- VOLUME_SIZE_VERY_LARGE_ARE_YOU_SURE
  The volume size is very large. It may not be possible to mirror this volume to older versions of the storage system. Are you sure?

Return codes

- BAD_REMOTE_VOLUME_NAME
  The secondary volume name does not exist.
- BAD_REMOTE_VOLUME_SIZE
  The primary and secondary volumes contain a different number of blocks.
- CONS_GROUP_BAD_NAME

The consistency group name does not exist.

- CONS_GROUP_HAS_MIRROR
  Mirroring is defined for this consistency group.
- CONS_GROUP_MIRRORING_NOT_SUPPORTED_IN_TARGET
  Consistency group mirroring is not supported by the target machine.
- LOCAL_MAX_HA_REACHED
  The maximum number of HyperSwap relationships is already reached on the local machine.
- NOT_ENOUGH_SPACE_ON_REMOTE_MACHINE
  Not enough free space to set the requested size of the secondary volume.
- VOLUME_BAD_NAME
  The volume name does not exist.
- VOLUME_IS_MASTER
  This local volume is already defined as a primary volume.
- VOLUME_IS_SLAVE
  The volume is defined as a secondary volume.
- VOLUME_HAS_OLVM
  An IBM Hyper-Scale Mobility relationship is defined for this volume.
- VOLUME_HAS_HA
  This operation is forbidden on a volume with a HyperSwap relation.
- TARGET_VOLUME_HAS_OLVM
  This target volume is part of an IBM Hyper-Scale Mobility relationship.
- TARGET_VOLUME_HAS_HA
  This operation is forbidden, if the target volume is a peer in a HyperSwap relation.
- REMOTE_VOLUME_EXISTS
  The secondary volume with the indicated name already exists. The name cannot be reused.
- REMOTE_MAX_VOLUMES_REACHED
  The maximum number of volumes on the remote machine is already reached.
- REMOTE_MAX_HA_REACHED
  The maximum number of HyperSwap relationships is already reached on the remote machine.
- REMOTE_CONS_GROUP_NOT_EMPTY
  The remote consistency group contains volumes.
- VOLUME_BAD_PREFIX
  The volume name has a reserved prefix.
- REMOTE_POOL_DOES_NOT_EXIST
  The pool does not exist on the remote machine.
- REMOTE_POOL_NOT_SPECIFIED
  Prior to creating a secondary volume, a pool must be defined on the remote machine.
- REMOTE_TARGET_NOT_CONNECTED
  There is currently no connection from the target system.
- VOLUME_IS_SNAPSHOT
  The operation is not permitted on snapshots.
- REMOTE_VOLUME_IS_SNAPSHOT
  The secondary volume is a snapshot.
- TARGET_BAD_NAME
  The target name does not exist.
- TARGET_BAD_TYPE
  The target machine is not an XIV machine.
- TARGET_NO_ACCESS
  No access permissions to the secondary machine.
- TARGET_NOT_CONNECTED
  There is currently no connection to the target system.
- REMOTE_VOLUME_LOCKED
  The secondary volume is locked.
- TIMEOUT
  A remote operation was not completed in time.
- VOLUME_HAS_MIRRORING_SNAPSHOTS

The volume has snapshots created by a previous mirroring process.

- SLAVE_VOLUME_NOT_FORMATTED
  The secondary volume is not formatted.
- TARGET_DOES_NOT_ACCEPT_XIV_COMMANDS
  The target system does not accept XIV management commands.
- REMOTE_CONS_GROUP_IS_MIRRORED
  Mirroring is defined for this remote consistency group.
- REMOTE_CONS_GROUP_BAD_NAME
  The remote consistency group name does not exist.
- REMOTE_VOLUME_IS_MASTER
  A volume on the remote machine is already defined as primary.
- REMOTE_VOLUME_IS_SLAVE
  A volume on the remote machine is already defined as secondary.
- REMOTE_MAX_MIRROR_CAPACITY_REACHED
  The maximum capacity for mirrored volumes is already reached on the remote machine.
- HA_RETRY_OPERATION
  An operation is in progress on this HyperSwap relation.

  Troubleshooting: Try issuing the command again in a few seconds.
- HA_INCOMPATIBLE_TARGET_VERSION
  The automatic failover is not supported between the system versions of the specified peers.
- NO_OFFLINE_INIT_TYPE_WITH_SLAVE_CREATION
  A new volume will be created as secondary. Offline initialization is meaningless.
- VOLUME_SIZE_ABOVE_LIMIT
  The specified volume size is above the limit.
- REMOTE_VOLUME_SIZE_ABOVE_LIMIT
  The specified volume size is above the limit of the remote machine.
- INVALID_SLICE_OFFSET
  Slice offset is illegal.
- VOLUME_IS_OLVM_PROXY
  The volume is in an IBM Hyper-Scale Mobility Proxy phase.
- REMOTE_VOLUME_IS_OLVM_PROXY
  The remote volume is in an IBM Hyper-Scale Mobility Proxy phase.
- ENCRYPTION_IN_PROGRESS
  The system is in the process of changing the encryption activation state.
- MIRROR_OF_SAME_TYPE_EXISTS_ON_VOLUME
  A mirror of this type is already defined on this volume.
- MIRROR_EXISTS_ON_TARGET
  The volume already has a mirror on this target.
- REMOTE_VOLUME_IS_MIRROR_MASTER
  The volume is primary in a mirror relationship, and cannot be secondary!
- MULTISITE_MAX_NUM_OF_MIRRORS_REACHED
  Failed to create the relation, the maximum allowed number of relations is already exceeded.
- REMOTE_VOLUME_TWO_SYNC_MIRRORS_NOT_ALLOWED
  Two synchronous mirrors were detected on the remote volume. This is not allowed.
- REMOTE_VOLUME_MIRROR_LOOP_DETECTED
  A mirror loop was detected on the remote volume. This means that there is a mirror on the remote system, whose target is this system. Therefore, you cannot create a mirror with this target on this system.
- DOMAIN_MAX_MIRRORS_REACHED
  The domain exceeds the maximum allowed number of mirrors.
- REMOTE_DOMAIN_MAX_VOLUMES_REACHED
  The maximum number of volumes in the remote machine domain is already reached.
- REMOTE_DOMAIN_HAS_NO_ACCESS_TO_TARGET
  The secondary machine domain has no access to the target.
- DOMAIN_HAS_NO_ACCESS_TO_TARGET
  The domain has no access to the target.
- REMOTE_DOMAIN_MAX_MIRRORS_REACHED

The maximum number of mirrors is already reached in the remote machine domain.

- DOMAIN_MAX_VOLUMES_REACHED
  The domain exceeds the maximum allowed number of volumes.
- REMOTE_VOLUME_HAS_DATA_MIGRATION
  Data migration is already defined for the secondary volume.
- REMOTE_VOLUME_MASTER_ASYNC_MIRROR_DETECTED
  An asynchronous primary mirror was detected on the remote volume. The operation not allowed.
- MAX_MULTISITES_REACHED
  The number of Multi-site objects exceeded the limit.
- REMOTE_VOLUME_HAS_MIRRORING_SNAPSHOTS
  The remote volume has snapshots created by a previous mirroring process.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support
- SYSTEM_OUT_OF_PHYSICAL_SPACE
  The operation not allowed while the system is out of physical space.
- REMOTE_DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier of the remote system is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support
- REMOTE_SYSTEM_OUT_OF_PHYSICAL_SPACE
  The operation not allowed while the remote system is out of physical space.
- SLAVE_VOLUME_IS_MAPPED
  The secondary volume is mapped.
- TARGET_HAS_NO_QUORUM_WITNESS
  The local target does not have a Quorum Witness defined.
- REMOTE_TARGET_HAS_NO_QUORUM_WITNESS
  The target on the remote system does not have a Quorum Witness defined.
- HA_TARGET_QUORUM_WITNESS_IS_NOT_ACTIVATED
  The Quorum Witness associated with the target is deactivated.
- HA_REMOTE_TARGET_QUORUM_WITNESS_IS_NOT_ACTIVATED
  The Quorum Witness associated with the remote target is deactivated.
- HA_CONNECTIVITY_NOT_SUFFICIENT
  The connectivity between the systems is not sufficient for the automatic failover.
- HA_LOCAL_PEER_HAS_NO_QUORUM_WITNESS_CONNECTIVITY
  The local peer connectivity to the Quorum Witness is not operational.
- HA_REMOTE_PEER_HAS_NO_QUORUM_WITNESS_CONNECTIVITY
  The remote peer connectivity to the Quorum Witness is not operational.
- MAX_NUM_OF_PROXY_VOLUME_REACHED
  Failed to create a mirror, because the maximum number of proxy volumes is exceeded.
- REMOTE_MAX_NUM_OF_PROXY_VOLUME_REACHED
  Failed to create a mirror, because the maximum number of remote proxy volumes is exceeded.
- REMOTE_MAX_METADATA_OBJECTS_REACHED
  The maximum number of metadata objects has been reached on a remote system.
- HA_PEER_QUORUM_WITNESS_CONFIGURATION_NOT_VERIFIED
  Connectivity between the peer and the Quorum Witness is not verified.
- HA_REMOTE_PEER_QUORUM_WITNESS_CONFIGURATION_NOT_VERIFIED
  Connectivity between the remote peer and the Quorum Witness is not verified.
- TARGET_PEER_NOT_HEALTHY
  The target peer is not identified as healthy.

  Troubleshooting: Check the Quorum Witness configuration.
- HOST_TYPE_IS_NOT_CONFIGURED
  Cannot associate a HyperSwap volume with a host of unconfigured type. IMPORTANT: Read the HyperSwap chapter in the 'Best Practices' document to understand the solution requirements.
- CONS_GROUP_IS_SECONDARY
  The consistency group is already defined as a secondary consistency group.

- CONS_GROUP_IS_PRIMARY
  The consistency group is already defined as a primary consistency group.
- RELATION_OF_SAME_TYPE_EXISTS_ON_CONS_GROUP
  A relation of this type is already defined on this consistency group.
- MULTISITE_INCOMPATIBLE_TARGET_VERSION
  Multi-site is not supported between the system versions of the specified peers.
- REMOTE_VOLUME_BAD_POOL
  The remote volume and remote consistency group belong to different storage pools.
- REMOTE_VOLUME_BELONGS_TO_CONS_GROUP
  The remote volume belongs to a consistency group.
- CONS_GROUP_HA_ROLE_MISMATCH
  All volumes in a HyperSwap consistency group must have the same mirroring role.
- HA_HIGH_AVAILABILITY_ENABLED_IN_VOL
  The consistency group's high availability is disabled but the volume's high availability is enabled.
- HA_HIGH_AVAILABILITY_DISABLED_IN_VOL
  The consistency group's high availability is enabled, but the volume's high availability is disabled.
- CONS_GROUP_RELATION_STATE_IS_NOT_INITIALIZING
  This operation is only allowed when the matching consistency group relation state is 'Initializing'.
- VOLUME_RELATION_IN_CG_CAN_NOT_BE_CREATED
  A second relation must first be added to the consistency group before it can be added to a member volume.
- CONS_GROUP_RELATION_PART_OF_MULTISITE_MISMATCH
  All volumes in a mirrored or HyperSwap consistency group must have the same part_of_multisite setting.
- CONS_GROUP_NOT_EMPTY
  This operation is only allowed on an empty consistency group.
- CONS_GROUP_HA_TARGET_MISMATCH
  All volumes in a mirrored consistency group must have the same HyperSwap target.
- CONS_GROUP_RELATION_INIT_TYPE_MISMATCH
  All volumes in a mirrored or HyperSwap consistency group must have the same init type setting.
- REMOTE_MULTISITE_CONS_GROUP_MEMBER_VOL_HAS_TOO_MANY_RELATIONS
  A remote volume belonging to the consistency group has too many relations.
- LOCAL_VOLUME_HAS_TOO_MANY_METADATA_OBJECTS
  Local volume has too many metadata objects.

  Troubleshooting: Contact IBM Support.

Viewing the status of HyperSwap volumes and consistency groups

Use the **ha_list** command to display the status of HyperSwap volumes and consistency groups.

```
ha_list [ < [ vol=VolName ] [ target=TargetName ] > | cg=cgName | scope=<cg|volume> ] [ domain=DomainName ]
```

Parameters

Name	Type	Description	Mandatory	Default
scope	Enumeration	List type: all HyperSwap mirrors, volumes, and consistency groups.	N	All (if no value is specified)
vol	Object name	Local volume name.	N	[none]
cg	Object name	Local consistency group name.	N	[none]
target	Object name	Remote target name.	N	[none]
domain	Object name	The domain name.	N	All Domains

The size, part, and time to synchronize are unknown if this is the Slave and connection is broken.

The following default parameters are shown:

- Name
- HA Object: Volume or CG
- Role: Master or Slave
- Remote System: The name of the target storage system
- Active: Yes or No
- Status: Initializing, Synchronized, or Unsynchronized
- Link Up: Yes or No
- Automatic Failover: The ability of a Slave volume to perform automatic failover: Active, Inactive, or N/A. Can be retrieved from either Master or Slave. N/A is returned for a volume that is part of a consistency group.

The following optional parameters can be listed by explicitly specifying the proper columns:

- Designation: Primary or Secondary
- Estimated Sync Time: Estimated time (in seconds) for synchronization to complete. This parameter is not available for a Slave volume or consistency group, if its Link Up status is No.
- Size To Synchronize: The amount of data (in MB) to synchronize. This parameter is not available for a Slave volume or consistency group, if its Link Up status is No
- Operational: Yes or No
- Sync Progress: The amount of synchronized data (in %)
- Automatic Failover Reason indicates why an automatic failover has been performed: N/A, User setting, Coordinated lapse, Unsynced, Quorum Witness issue
- IO Service indicates the local peer ability to serve I/O: Active or Unavailable

Field ID	Field output	Description	Default position
local_peer_name	Name	N/A	1
ha_object	HA Object	N/A	2
current_role	Role	N/A	3
target_name	Remote System	N/A	4
active	Active	N/A	5
sync_state	Status	N/A	6
ha_connected	Link Up	N/A	7
automatic_failover	Automatic Failover	N/A	8
automatic_failover_reason	Automatic Failover Reason	N/A	N/A
remote_peer_name	Remote Peer	N/A	N/A
designation	Designation	N/A	N/A
size_to_synchronize	Size To Sync (MiB)	N/A	N/A
operational	Operational	N/A	N/A
sync_progress	Sync Progress (%)	N/A	N/A
mirror_error	Mirror Error	No Error, Secondary pool exhausted, Configuration error or No thin provisioning resources	N/A
crash_consistent	Crash Consistency	N/A	N/A
validate	Validation	N/A	N/A
ha_high_availability_state	HA High Availability State	N/A	N/A
ha_unavailable_reason	HA Unavailable Reason	N/A	N/A
ha_sync_state	HA Sync State	N/A	N/A
ha_object_state	HA object State	N/A	N/A
io_service	I/O Service	N/A	N/A
part_of_multisite	part_of_multisite	N/A	N/A
arch	Remote Arch	N/A	N/A
multisite_uid	multisite_uid	N/A	N/A

Example:

```
ha_list
```

Output:

```
Name   HA Object   Role     Remote System          Active   Status           Link Up   Automatic Failover
------ ----------- -------- ---------------------- -------- ---------------- --------- --------------------
ha_1   Volume      Master   target-3726085-0008    yes      Synchronized     yes       Active
ha_2   Volume      Slave    target-3726085-0008    yes      Synchronized     yes       Active
ha_3   Volume      Slave    target-3726085-0008    yes      Initializing     yes       Inactive
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Activating a HyperSwap relation

Use the **ha_activate** command to activate a HyperSwap relation of volumes or consistency groups.

```
ha_activate < vol=VolName | cg=cgName > [ target=TargetName ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	Master volume name. If this parameter is not used, the cg parameter must be specified.	N	N/A
cg	Object name	Master consistency group name. If this parameter is not used, the vol parameter must be specified.	N	N/A
target	Object name	Target HyperSwap relation name.	N	[none]

The command updates the Quorum Witness about the state of the HyperSwap relation after the operation. If high availability is enabled, the command will fail if the Quorum Witness update fails. If high availability is disabled, the command may succeed even if the Quorum Witness update fails.

The following is required for a successful command completion:

- The Master and Slave are configured on the same Quorum Witness
- The connectivity of Master and Slave to the Quorum Witness is healthy
- The specified object is a Master
- The connectivity between Master and Slave is sufficient for HyperSwap (that is, the link_up attribute of the **ha_list** command output is *yes*

If the relation is already in the Active state, nothing is done and a success code is returned.

A HyperSwap relation cannot be activated, if:

- The command is issued on a Master that did not receive acknowledgment from the Slave following the **cg_add_vol** or **cg_remove_vol** command, due to the command's timeout or to an unexpected failure. In this case, the command fails and the HA_CONFIGURATION_ERROR code is returned. This means that the member lists of the mirror consistency group peers are not the same.
- The command is issued on a Master that did not receive acknowledgment from the Slave following a **vol_resize** command, due to the command's timeout or to an unexpected failure. In this case, the command fails and the HA_CONFIGURATION_ERROR code is returned. This means that the sizes of the HyperSwap relation volume peers are not the same.

Example:

```
ha_activate target=123 vol=regular_volume
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.
- CONS_GROUP_BAD_NAME
  The consistency group name does not exist.
- LOCAL_PEER_IS_NOT_MASTER
  The local peer is not primary.
- HA_CONFIGURATION_ERROR
  The HyperSwap relation's local configuration does not match its remote configuration.
- REMOTE_MAX_VOLUMES_REACHED
  The maximum number of volumes on the remote machine is already reached.
- SYNC_ALREADY_ACTIVE
  Synchronization is already active.
- VOLUME_BELONGS_TO_HA_CONS_GROUP
  This volume belongs to a HyperSwap consistency group. The operation can be carried out on a volume which is not part of a consistency group, or on the consistency group itself.
- DOMAIN_MAX_VOLUMES_REACHED
  The domain exceeds the maximum allowed number of volumes.
- HA_CAN_NOT_BE_ACTIVATED
  The automatic failover cannot be activated.
- HA_CONS_GROUP_MEMBERSHIP_MISMATCH

The HyperSwap consistency group contains different primary and secondary volumes. This may have happened because the cg_add_vol or cg_remove_vol command was previously issued, but the primary did not receive an acknowledgment from the secondary until the command timed out, or due to any other unexpected failure.

- HA_SIZE_MISMATCH
  The sizes of the primary and secondary volumes in this HyperSwap relation are different.
- HA_RETRY_OPERATION
  An operation is in progress on this HyperSwap relation.

  Troubleshooting: Try issuing the command again in a few seconds.
- TARGET_BAD_NAME
  The target name does not exist.
- VOLUME_TARGET_MISMATCH
  The volume and target do not match.
- CONS_GROUP_BAD_TARGET
  The target name does not match the consistency group.
- VOLUME_TOO_MANY_ACTIVE_MIRRORS
  This command cannot be issued if more than one mirror is active on the volume.
- REMOTE_DOMAIN_MAX_VOLUMES_REACHED
  The maximum number of volumes in the remote machine domain is already reached.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support
- REMOTE_DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier of the remote system is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support
- REMOTE_SYSTEM_OUT_OF_PHYSICAL_SPACE
  The operation not allowed while the remote system is out of physical space.
- TARGET_PEER_NOT_HEALTHY
  The target peer is not identified as healthy.

  Troubleshooting: Check the Quorum Witness configuration.
- TARGET_HAS_NO_QUORUM_WITNESS
  The local target does not have a Quorum Witness defined.
- REMOTE_TARGET_HAS_NO_QUORUM_WITNESS
  The target on the remote system does not have a Quorum Witness defined.
- VOLUME_IS_NOT_HA
  This is not a HyperSwap volume.
- CONS_GROUP_IS_NOT_HA
  The local consistency group does not have HyperSwap definitions.
- HA_CONNECTIVITY_NOT_SUFFICIENT
  The connectivity between the systems is not sufficient for the automatic failover.
- HA_PEER_QUORUM_WITNESS_CONFIGURATION_NOT_VERIFIED
  Connectivity between the peer and the Quorum Witness is not verified.
- HA_REMOTE_PEER_QUORUM_WITNESS_CONFIGURATION_NOT_VERIFIED
  Connectivity between the remote peer and the Quorum Witness is not verified.
- HA_TARGET_QUORUM_WITNESS_IS_NOT_ACTIVATED
  The Quorum Witness associated with the target is deactivated.
- HA_REMOTE_TARGET_QUORUM_WITNESS_IS_NOT_ACTIVATED
  The Quorum Witness associated with the remote target is deactivated.
- CONS_GROUP_MEMBER_VOL_IS_MISSING_A_RELATION
  A volume in the consistency group is missing the volume-level Multi-site/mirror/HyperSwap relation.
- REMOTE_MULTISITE_IS_NOT_SMASTER
  The remote peer is not the SMaster.
- MULTISITE_SMASTER_INVALID_CONFIGURATION
  The Multi-site SMaster-Master relation configuration is invalid.
- MULTISITE_BAD_GLOBAL_ID
  The Multi-site global ID does not exist.

Deactivating a HyperSwap relation

Use the **ha_deactivate** command to deactivate HyperSwap volumes or consistency groups.

```
ha_deactivate  < vol=<vol1[,vol2]...> | cg=cgName > [ target=TargetName ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	Master volume name or a list of master volumes. If this paramater is not used, the cg parameter must be specified.	N	N/A
cg	Object name	Master consistency group name or a list of master consistency groups. If this parameter is not used, the vol parameter must be specified.	N	N/A
target	Object name	Target HyperSwap relation name.	N	[none]

This command deactivates a HyperSwap relation and changes its status to *Inactive*. While in the *Inactive* state, only the Master volume is updated, as opposed to the *Active* state, where the Slave volume is updated together with the Master volume.

Example:

```
ha_deactivate vol=regular_volume -y
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.
- VOLUME_IS_NOT_HA
  This is not a HyperSwap volume.
- CONS_GROUP_BAD_NAME
  The consistency group name does not exist.
- CONS_GROUP_IS_NOT_HA
  The local consistency group does not have HyperSwap definitions.
- LOCAL_PEER_IS_NOT_MASTER
  The local peer is not primary.
- SYNC_ALREADY_INACTIVE
  Synchronization is already inactive.
- VOLUME_BELONGS_TO_HA_CONS_GROUP
  This volume belongs to a HyperSwap consistency group. The operation can be carried out on a volume which is not part of a consistency group, or on the consistency group itself.
- HA_RETRY_OPERATION
  An operation is in progress on this HyperSwap relation.

  Troubleshooting: Try issuing the command again in a few seconds.
- TARGET_BAD_NAME
  The target name does not exist.
- VOLUME_TARGET_MISMATCH
  The volume and target do not match.
- CONS_GROUP_BAD_TARGET
  The target name does not match the consistency group.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

Troubleshooting: Contact IBM Support

Deleting a HyperSwap relation

Use the **ha_delete** command to delete a HyperSwap relation.

```
ha_delete < vol=VolName | cg=cgName > [ target=TargetName ] [ force_on_slave=<Yes|No> ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	Local Master volume name. If this parameter is not used, the cg parameter must be specified.	N	N/A
cg	Object name	Local Master consistency group name. If this parameter is not used, the vol parameter must be specified.	N	N/A
target	Object name	Target HyperSwap relation name.	N	[none]
force_on_slave	Boolean	Forces the deletion of the HyperSwap relation on the Slave peer. Available only if the Slave is in the Initializing state.	N	no

The command can only be issued on the Master. It deletes only the definition of the HyperSwap relation. Neither the volumes themselves, nor their snapshots, are deleted.

To delete a HyperSwap relation, the following preconditions must be met:

- The HyperSwap relation is in the Standby state, in other words, it has just been created or deactivated.
- The communication is established. If there is no communication, the HyperSwap relation is only deleted on the Master. When the communication resumes, a configuration error will be issued on the Slave.
- The Slave volume is not mapped.

After deleting a HyperSwap relation, both its peers are labeled as none, meaning that they are no longer configured as either Master or Slave.

After the command's successful completion, the Slave volume's SCSI identification is replaced. The volume name, external ID, lock state, and metadata remain identical to the Master volume values.

Deletion when the HyperSwap relation is inactive or when the connectivity has failed

On an active Master

On an active Master volume, the command can be executed as described above.

On an active Slave

Important:
Never map the Slave volume to a host. If you need to read or update data, use the **vol_copy** command to copy the data to a new volume, and map this new volume to the host.

Deleting a HyperSwap relation when the communication between the peers is down, deletes only the HyperSwap relation from the Master. To delete the HyperSwap relation from the Slave:

- Run the **ha_change_role** command to turn the Slave into the Master.
- Run **ha_delete**.

The force_on_slave parameter can be used only if the HyperSwap relation is in the Initializing phase. In any other state, change the Slave into the Master, and run **ha_delete**.

On a disconnected Master or Slave

To delete the HyperSwap relation on a disconnected peer volume, re-connect the peer volume and run **ha_delete**.

If re-connecting the peer volume is not possible, contact IBM support.

Example:

```
ha_delete vol=regular_volume -y
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed

User Category	Permission
Technicians	Disallowed

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_DELETE_CG_HA
  Are you sure you want to delete the HyperSwap relations of the consistency group and of all the volumes in it?

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.
- VOLUME_IS_NOT_HA
  This is not a HyperSwap volume.
- CONS_GROUP_IS_NOT_HA
  The local consistency group does not have HyperSwap definitions.
- CONS_GROUP_BAD_NAME
  The consistency group name does not exist.
- LOCAL_PEER_IS_NOT_MASTER
  The local peer is not primary.
- HA_IS_ACTIVE
  HyperSwap is currently active.
- FORCE_DELETE_NOT_ALLOWED_ON_MASTER
  Deletion needs to be forced on secondary mirrors only.
- VOLUME_BELONGS_TO_HA_CONS_GROUP
  This volume belongs to a HyperSwap consistency group. The operation can be carried out on a volume which is not part of a consistency group, or on the consistency group itself.
- HA_RETRY_OPERATION
  An operation is in progress on this HyperSwap relation.

  Troubleshooting: Try issuing the command again in a few seconds.
- HA_IS_NOT_INITIALIZING
  Deleting a HyperSwap relation is permitted only during the initialization phase.
- TARGET_BAD_NAME
  The target name does not exist.
- VOLUME_TARGET_MISMATCH
  The volume and target do not match.
- CONS_GROUP_BAD_TARGET
  The target name does not match the consistency group.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support
- SLAVE_VOLUME_IS_MAPPED
  The secondary volume is mapped.
- MIRROR_CONS_GROUP_MEMBERSHIP_MISMATCH
  The mirrored consistency group contains different volumes on the primary and secondary machines. This problem occurs whenever the cg_add_vol or cg_remove_vol commands were previously issued, and the primary machine did not receive an acknowledgment from the secondary machine until the command timed out, or due to any other unexpected failure.
- VOLUME_IS_AN_UNAVAILABLE_HYPERSWAP_PEER
  The operation is not permitted on a HyperSwap target which is unavailable for I/O.
- HA_ASSOCIATED_WITH_MULTISITE
  This HyperSwap relation is associated with a defined Multi-site, the operation is not allowed.
- HA_PART_OF_MULTISITE
  This HyperSwap relation is part of a Multi-site relation.

Switching roles between Master and Slave volumes

Use the **ha_switch_roles** command to switch roles between Master and Slave volumes.

```
ha_switch_roles <vol=VolName | cg=cgName> [ target=TargetName ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	Local volume name. If this parameter is not used, the cg parameter must be specified.	N	N/A
cg	Object name	Local consistency group name. If this parameter is not used, the vol parameter must be specified.	N	N/A
target	Object name	N/A	N	[none]

The command can only be issued on the Master volume, and only if the HyperSwap relation is activated and synchronized.

When this command is issued, the system performs all pending writes, and only after all pending writes have been committed, the roles are switched.

Following the successful completion of the command:

- The volume that was previously the Master becomes the Slave
- The volume that was previously the Slave becomes the Master
- The HyperSwap relation remains active

Example:

```
ha_switch_roles vol=DBVolume target=SecondarySite
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_SWITCH_ROLES
  Are you sure you want to switch the roles in this relation?
- ARE_YOU_SURE_YOU_WANT_TO_SWITCH_THE_PEER_ROLES
  The system that owns the primary volume is currently not connected to the Quorum Witness. If you switch the peer roles, the automatic failover may become unavailable. Are you sure you want to continue?

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.
- CONS_GROUP_BAD_NAME
  The consistency group name does not exist.
- LOCAL_PEER_IS_NOT_MASTER
  The local peer is not primary.
- HA_IS_NOT_SYNCHRONIZED
  The HyperSwap relation is not synchronized.
- VOLUME_HAS_DATA_MIGRATION
  Data Migration is defined for this volume.
- REMOTE_TARGET_NOT_CONNECTED
  There is currently no connection from the target system.
- VOLUME_BELONGS_TO_HA_CONS_GROUP
  This volume belongs to a HyperSwap consistency group. The operation can be carried out on a volume which is not part of a consistency group, or on the consistency group itself.
- HA_CONNECTIVITY_NOT_SUFFICIENT
  The connectivity between the systems is not sufficient for the automatic failover.
- HA_RELATION_MASTER_COULD_NOT_UPDATE_QW_AFTER_RETURN_TO_GOOD_STATE
  The primary volume had control over the relationship during a past failure, and was unable to update the Quorum Witness after the recovery.

- HA_HAS_SYNC_JOB
  This operation is not permitted on a HyperSwap relation with active sync jobs.
- HA_RETRY_OPERATION
  An operation is in progress on this HyperSwap relation.

  Troubleshooting: Try issuing the command again in a few seconds.
- REMOTE_HA_IS_NOT_ACTIVE
  The remote peer in this HyperSwap relation is not active.
- TARGET_BAD_NAME
  The target name does not exist.
- VOLUME_TARGET_MISMATCH
  The volume and target do not match.
- CONS_GROUP_BAD_TARGET
  The target name does not match the consistency group.
- COMMAND_NOT_SUPPORTED_FOR_OLVM_VOLUMES
  This command is not supported for IBM Hyper-Scale Mobility volumes.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support
- SYSTEM_OUT_OF_PHYSICAL_SPACE
  The operation not allowed while the system is out of physical space.
- VOLUME_IS_NOT_HA
  This is not a HyperSwap volume.
- CONS_GROUP_IS_NOT_HA
  The local consistency group does not have HyperSwap definitions.
- REMOTE_MAY_NOT_HAVE_COMPLETED_THE_OPERATION
  The operation may be not yet completed on the remote target.
- VOLUME_HAS_MULTIPLE_RELATIONS
  The volume has multiple relations. The operation is not allowed or a target must be specified.
- CONS_GROUP_HAS_MULTIPLE_RELATIONS
  The consistency group has multiple relations. The operation is not allowed or a target must be specified.

Changing a peer role in a HyperSwap volume

Use the **ha_change_role** command to change the role of a local HyperSwap relation peer from Master to Slave or from Slave to Master.

```
ha_change_role <vol=VolName | cg=cgName> [ target=TargetName ] [ new_role=<Master|Slave|None> ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	Local volume name. Must be specified if the command is applied to a volume.	N	N/A
cg	Object name	Consistency group name. Must be specified if the command is applied to a consistency group.	N	N/A
target	Object name	Target HyperSwap relation name.	N	[none]
new_role	Enumeration	Role name of the peer. If not specified, the command swaps peer roles between Master and Slave.	N	[none]

This command changes the role of the local peer from Master to Slave or from Slave to Master when the HyperSwap relation is not activated. The command should be issued on both peers before the relation is activated again, so that upon reconnection there still will be one Master and one Slave.

The command is used during recovery after an automatic failover, or in order to perform a manual failover when the automatic failover did not take place.

For a successful role change from Master to Slave, the volume can be in any phase, except *Initializing*. The Master ceases serving host requests, and is set to accept replication from the other peer as a Slave.

To successfully change a Slave to a Master, the HyperSwap relation must be deactivated.

Before changing a Slave to a Master:

- Make sure that the original Master is not available and cannot become available while the other peer is a Master. To verify this, run the **ha_list** command on the Master, and check the value of the attribute IO Service. Only if the returned value is Unavailable, proceed with issuing the **ha_change_role** command on the Slave.

- Stop the applications using the HyperSwap volume(s). Note that each application must be stopped completely and not merely paused in order to make sure that it does not use any cached state when accessing the volume for the first time after the role change.

Warning: Failure to fulfill both of the above requirements may result in a data integrity issue.

After a Slave is successfully changed to a Master, the volume starts accepting requests from hosts. Upon explicit activation, it starts replicating to the other peer (the original Master).

If the synchronous mirroring is interrupted in the middle of the re-synchronization process, the Slave volume may very probably be inconsistent. The last consistent image of the Slave volume is preserved in the last_consistent snapshot (LCS), which is automatically created immediately before the re-synchronization starts. If the LCS exists, the command emits a warning: Are you sure you want the mirror/HyperSwap local peer to become primary? The local peer has a last-consistent snapshot. In this case, the administrator must choose whether to use the existing contents of the previous Slave volume, which may be inconsistent, or revert the previous Slave volume to its last_consistent snapshot before issuing the **ha_change_role** command.

Example:

```
ha_change_role vol=regular_volume
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- SOME_DATA_WILL_BE_LOST_ARE_YOU_SURE
  Are you sure you want the mirror/HyperSwap local peer to become secondary and lose the data that was not replicated?
- ARE_YOU_SURE_YOU_WANT_TO_CHANGE_A_PEER_WITH_LCS_TO_MASTER
  Are you sure you want the mirror/HyperSwap local peer to become primary? The local peer has a last-consistent snapshot.

Return codes

- VOLUME_IS_NOT_HA
  This is not a HyperSwap volume.
- CONS_GROUP_IS_NOT_HA
  The local consistency group does not have HyperSwap definitions.
- HA_IS_ACTIVE
  HyperSwap is currently active.
- HA_IS_INITIAL
  The operation is not permitted during the HyperSwap relation initialization phase.
- VOLUME_BELONGS_TO_HA_CONS_GROUP
  This volume belongs to a HyperSwap consistency group. The operation can be carried out on a volume which is not part of a consistency group, or on the consistency group itself.
- VOLUME_BAD_NAME
  The volume name does not exist.
- CONS_GROUP_BAD_NAME
  The consistency group name does not exist.
- VOLUME_HAS_DATA_MIGRATION
  Data Migration is defined for this volume.
- HA_RETRY_OPERATION
  An operation is in progress on this HyperSwap relation.

  Troubleshooting: Try issuing the command again in a few seconds.
- HA_HAS_NO_SYNCHED_SNAPSHOT
  This HyperSwap volume does not have a synchronized snapshot.
- MASTER_CANNOT_BE_DEMOTED
  The primary volume cannot be demoted to secondary. Peer status mismatch.

- TARGET_BAD_NAME
  The target name does not exist.
- VOLUME_TARGET_MISMATCH
  The volume and target do not match.
- CONS_GROUP_BAD_TARGET
  The target name does not match the consistency group.
- COMMAND_NOT_SUPPORTED_FOR_OLVM_VOLUMES
  This command is not supported for IBM Hyper-Scale Mobility volumes.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support
- SYSTEM_OUT_OF_PHYSICAL_SPACE
  The operation not allowed while the system is out of physical space.
- VOLUME_HAS_MULTIPLE_RELATIONS
  The volume has multiple relations. The operation is not allowed or a target must be specified.
- CONS_GROUP_HAS_MULTIPLE_RELATIONS
  The consistency group has multiple relations. The operation is not allowed or a target must be specified.
- HA_PART_OF_MULTISITE
  This HyperSwap relation is part of a Multi-site relation.

Restoring the availability of a Master volume

Use the **ha_restore_availability** command to restore the availability of a Master volume, that became unavailable due to a failure.

```
ha_restore_availability <vol=VolName | cg=cgName>
```

Parameters

Name	Type	Description	Mandatory
vol	Object name	Local volume name. Must be specified if the command is applied to a volume.	N
cg	Object name	CG name Must be specified if the command is applied to a consistency group.	N

As a result of some failure scenarios, the Master may assume that the Slave has performed an automatic failover, and stop handling I/O. In such cases, the user can choose to disable the remote peer and manually restore the availability of the Master.

Prior to issuing this command, make sure that the remote target is a Slave. If it is a Master, make sure that it is not available. To verify this, run the **ha_list** command on the remote target, and check the value of the attribute *IO Service*. Only if the returned value is *Unavailable*, proceed with issuing the **ha_restore_availability** command.

Upon issuing this command, the following occurs:

- An event is generated
- The Master volume becomes available

Example:

```
ha_restore_availability vol=regular_volume
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- VOLUME_IS_NOT_HA

This is not a HyperSwap volume.

- CONS_GROUP_IS_NOT_HA
  The local consistency group does not have HyperSwap definitions.
- CONS_GROUP_BAD_NAME
  The consistency group name does not exist.
- HA_IS_ACTIVE
  HyperSwap is currently active.
- VOLUME_BELONGS_TO_HA_CONS_GROUP
  This volume belongs to a HyperSwap consistency group. The operation can be carried out on a volume which is not part of a consistency group, or on the consistency group itself.
- VOLUME_BAD_NAME
  The volume name does not exist.
- LOCAL_PEER_IS_NOT_MASTER
  The local peer is not primary.

Creating a HyperSwap volume snapshot (ad hoc sync job)

Use the **ha_create_snapshot** command to simultaneously create snapshots on both peers of a HyperSwap relation.

```
ha_create_snapshot <vol=VolName | cg=cgName> [ target=TargetName ] name=Name  [ slave_name=SnapshotName ]
 [ delete_priority=del_value ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	The name of the volume to create a snapshot for.	N	N/A
cg	Object name	Local master consistency group name.	N	N/A
target	Object name	Target HyperSwap relation name. Mandatory if 2 mirrors are defined on the volume.	N	[none]
name	Object name	The name of the new snapshot.	Y	N/A
slave_name	Object name	The name of the new snapshot on the slave.	N	[none]
delete_priority	Integer	The deletion priority of the volume's snapshot.	N	1

This command takes a snapshot of the source peer (Master) and the target peer (Slave) at exactly the same time. The snapshots created concurrently on the Master and Slave are identical.

Prerequisite:

- The HyperSwap relation is operational and synchronized.

The snapshots created by this command can be managed with regular snapshot commands. For example, to delete these snapshots, issue the **snapshot_delete** command at each peer.

Example:

```
ha_create_snapshot vol=simpleVolume target=myTarget name=mySnapshot slave_name=slave1 delete_priority=3
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission	Condition
Storage administrator	Allowed	N/A
Storage integration administrator	Allowed	N/A
Application administrator	Conditionally Allowed	The volume is mapped to a host or a cluster associated with the user. If a snapshot overwrite is used, the target snapshot must be one created by a server administrator.
Security administrator	Disallowed	N/A
Read-only users	Disallowed	N/A
Technicians	Disallowed	N/A

Return codes

- CONS_GROUP_MISMATCH
  The snapshot group does not match the consistency group volumes.
- CONS_GROUP_EMPTY
  The operation is not allowed on an empty consistency group.

- CONS_GROUP_BAD_NAME
  The consistency group name does not exist.
- VOLUME_IS_NOT_HA
  This is not a HyperSwap volume.
- CONS_GROUP_IS_NOT_HA
  The local consistency group does not have HyperSwap definitions.
- LOCAL_PEER_IS_NOT_MASTER
  The local peer is not primary.
- HA_IS_NOT_SYNCHRONIZED
  The HyperSwap relation is not synchronized.
- HA_RETRY_OPERATION
  An operation is in progress on this HyperSwap relation.

  Troubleshooting: Try issuing the command again in a few seconds.
- HA_IS_NON_OPERATIONAL
  This HyperSwap volume is not operational.
- MAX_VOLUMES_REACHED
  The maximum allowed number of volumes is already reached.
- DOMAIN_MAX_VOLUMES_REACHED
  The domain exceeds the maximum allowed number of volumes.
- OPERATION_NOT_ALLOWED_ON_LOOPBACK
  The requested operation is not allowed on a loopback target.
- OVERWRITE_SNAPSHOT_BAD_NAME
  The snapshot name does not exist.
- OVERWRITE_SNAPSHOT_GROUP_DOES_NOT_BELONG_TO_GIVEN_GROUP
  The snapshot group belongs to another consistency group.
- POOL_SNAPSHOT_LIMIT_REACHED
  There is not enough space to create a snapshot.
- REMOTE_POOL_SNAPSHOT_LIMIT_REACHED
  There is not enough space on the remote target for creating a snapshot.
- REMOTE_MAX_VOLUMES_REACHED
  The maximum number of volumes on the remote machine is already reached.
- REMOTE_MAX_SNAPSHOTS_FOR_VOLUME_REACHED
  The maximum allowed number of snapshots per volume is already reached on a remote machine whose version is not 10.2.4.
- REMOTE_VOLUME_IS_MASTER
  A volume on the remote machine is already defined as primary.
- REMOTE_VOLUME_IS_SNAPSHOT
  The secondary volume is a snapshot.
- REMOTE_SNAPSHOT_NAME_EXISTS
  The remote snapshot name already exists.
- REMOTE_SNAPSHOT_ILLEGAL_PRIORITY
  Illegal snapshot priority (remote); must be an integer between 1 and 4.
- REMOTE_SNAPSHOT_GROUP_NAME_EXISTS
  The remote snapshot group name already exists.
- REMOTE_SNAPSHOT_GROUP_ILLEGAL_PRIORITY
  Illegal snapshot group priority (remote); must be an integer between 1 and 4.
- REMOTE_SNAPSHOT_GROUP_BAD_PREFIX
  The remote snapshot group name has a reserved prefix.
- REMOTE_SNAPSHOT_BAD_PREFIX
  The remote snapshot name has a reserved prefix.
- REMOTE_CONS_GROUP_EMPTY
  The operation is not allowed on an empty consistency group (remote).
- REMOTE_CONS_GROUP_MISMATCH
  The remote snapshot group does not match the consistency group volumes.
- SNAPSHOT_ILLEGAL_PRIORITY
  Illegal snapshot priority; must be an integer between 1 and 4.
- SNAPSHOT_IS_INTERNAL

Internal snapshots cannot be mapped, modified or deleted.

- SNAPSHOT_GROUP_IS_INTERNAL
  Internal snapshots cannot be mapped, modified, or deleted.
- SNAPSHOT_GROUP_NAME_EXISTS
  The snapshot group name already exists.
- SNAPSHOT_GROUP_ILLEGAL_PRIORITY
  Illegal snapshot group priority; must be an integer between 1 and 4.
- SNAPSHOT_GROUP_BAD_NAME
  The snapshot group name does not exist.
- SNAPSHOT_GROUP_BAD_PREFIX
  The snapshot group name has a reserved prefix.
- SNAPSHOT_IS_PART_OF_SNAPSHOT_GROUP
  The snapshot is part of a snapshot group.
- VOLUME_BAD_PREFIX
  The volume name has a reserved prefix.
- VOLUME_BELONGS_TO_HA_CONS_GROUP
  This volume belongs to a HyperSwap consistency group. The operation can be carried out on a volume which is not part of a consistency group, or on the consistency group itself.
- VOLUME_EXISTS
  The volume name already exists.
- VOLUME_BAD_NAME
  The volume name does not exist.
- VOLUME_IS_NOT_CONSISTENT_SLAVE
  The operation not allowed on an inconsistent secondary volume.
- VOLUME_IS_SNAPSHOT
  The operation is not permitted on snapshots.
- CONS_GROUP_BAD_TARGET
  The target name does not match the consistency group.
- TARGET_BAD_NAME
  The target name does not exist.
- VOLUME_TARGET_MISMATCH
  The volume and target do not match.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support
- MAX_SNAPSHOTS_PER_VOLUME_REACHED
  The maximum allowed number of snapshots is already reached.
- REMOTE_MAX_SNAPSHOTS_PER_VOLUME_REACHED
  The maximum allowed number of snapshots is already reached on the remote ssytem.
- REMOTE_DOMAIN_MAX_VOLUMES_REACHED
  The maximum number of volumes in the remote machine domain is already reached.
- REMOTE_CONS_GROUP_BAD_NAME
  The remote consistency group name does not exist.
- SNAPSHOT_CAN_NOT_BE_CREATED_REMOTE_CONS_GROUP_IO_IS_NOT_PAUSED
  The snapshot group will not be created since the remote consistency group is not in a stopped state.
- SNAPSHOT_CAN_NOT_BE_CREATED_REMOTE_CONS_GROUP_DEFINITION_CHANGED
  The snapshot group will not be created since the volumes in the remote consistency group have changed since the io_pause command was issued.
- REMOTE_OVERWRITE_SNAPSHOT_GROUP_DOES_NOT_BELONG_TO_GIVEN_GROUP
  The remote snapshot group belongs to another consistency group.

Changing the designation of HyperSwap relation peers

Use the **ha_change_designation** command to change the designation of HyperSwap relation peers from Primary to Secondary and vice versa.

```
ha_change_designation < vol=VolName | cg=cgName > [ target=TargetName ]
 [ new_designation=<Primary|Secondary|None> ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	Master volume name.	N	N/A
cg	Object name	Master consistency group name.	N	N/A
target	Object name	The name of the target.	N	[none]
new_designation	Enumeration	The new designation of the peer If not specified, the command swaps the designation of the Primary and Secondary peers.	N	none

The designation in a HyperSwap relation reflects the user's decision where the Primary (Master) and the Secondary (Slave) peers should be located. The actual roles performed by the two peers at any given moment may differ from their designations, as a result of a manual role change or an automatic failover.

This command is issued on the Primary peer and affects both peers. For the command to be successfully completed, the HyperSwap relation has to be operational.

Specifying the new designations is not mandatory. If they are not specified, the command swaps the designations of both peers: the Primary changes to Secondary, and the Secondary changes to Primary.

Example:

```
ha_change_designation vol=regular_volume new_designation=Secondary
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.
- VOLUME_IS_NOT_HA
  This is not a HyperSwap volume.
- CONS_GROUP_BAD_NAME
  The consistency group name does not exist.
- CONS_GROUP_IS_NOT_HA
  The local consistency group does not have HyperSwap definitions.
- LOCAL_PEER_IS_NOT_MASTER
  The local peer is not primary.
- MIRROR_DESIGNATION_NOT_SUPPORTED_BY_TARGET
  The mirror's target does not support mirror role designation.
- HA_IS_NON_OPERATIONAL
  This HyperSwap volume is not operational.
- VOLUME_BELONGS_TO_HA_CONS_GROUP
  This volume belongs to a HyperSwap consistency group. The operation can be carried out on a volume which is not part of a consistency group, or on the consistency group itself.
- TARGET_BAD_NAME
  The target name does not exist.
- VOLUME_TARGET_MISMATCH
  The volume and target do not match.
- CONS_GROUP_BAD_TARGET
  The target name does not match the consistency group.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support

Enabling high availability of a HyperSwap relation

Use the **ha_high_availability_enable** command to enable high availability (automatic failover) of a HyperSwap relation.

```
ha_high_availability_enable < vol=VolName | cg=cgName >
```

Parameters

Name	Type	Description	Mandatory
vol	Object name	Master volume.	N
cg	Object name	Master consistency group name or a list of master consistency groups.	N

The high availability is enabled in order to allow automatic failover, after the high availability had been disabled in the past. The actual state of the automatic failover, as reported in **ha_list**, depends on several factors, including the user enablement and data synchronization state. It is possible that even though the high availability is enabled by this command, the automatic failover remains inactive due to other factors.

The command requires that the HyperSwap relation be active. It updates both peers, but must be issued only on the Master peer volume or consistency group.

Example:

```
ha_high_availability_enable vol=regular_volume
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.
- CONS_GROUP_BAD_NAME
  The consistency group name does not exist.
- VOLUME_IS_NOT_HA
  This is not a HyperSwap volume.
- CONS_GROUP_IS_NOT_HA
  The local consistency group does not have HyperSwap definitions.
- TARGET_NOT_CONNECTED
  There is currently no connection to the target system.
- REMOTE_TARGET_NOT_CONNECTED
  There is currently no connection from the target system.
- LOCAL_PEER_IS_NOT_MASTER
  The local peer is not primary.
- VOLUME_BELONGS_TO_HA_CONS_GROUP
  This volume belongs to a HyperSwap consistency group. The operation can be carried out on a volume which is not part of a consistency group, or on the consistency group itself.
- HA_IS_NOT_OPERATIONAL
  This HyperSwap relation is not operational. The operation cannot be carried out on a non-operational HyperSwap relation.
- HA_HIGH_AVAILABILITY_IS_ALREADY_ENABLED
  Automatic failover is already enabled (valid only for ha_high_availability_enable).
- HA_PEER_QUORUM_WITNESS_CONFIGURATION_NOT_VERIFIED
  Connectivity between the peer and the Quorum Witness is not verified.
- HA_REMOTE_PEER_QUORUM_WITNESS_CONFIGURATION_NOT_VERIFIED
  Connectivity between the remote peer and the Quorum Witness is not verified.

Disabling high availability of a HyperSwap relation

Use the **ha_high_availability_disable** command to disable the high availability (automatic failover) of a HyperSwap relation.

```
ha_high_availability_disable < vol=VolName | cg=cgName >
```

Parameters

Name	Type	Description	Mandatory
vol	Object name	Master volume.	N
cg	Object name	Master consistency group name or a list of master consistency groups.	N

The user may choose to disable the high availability and thereby prevent automatic failover in some maintenance scenarios, notably when the Quorum Witness is being replaced.

The command updates both peers, but must be issued only on the Master peer volume or consistency group.

Example:

```
ha_high_availability_disable vol=regular_volume
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.
- CONS_GROUP_BAD_NAME
  The consistency group name does not exist.
- VOLUME_IS_NOT_HA
  This is not a HyperSwap volume.
- CONS_GROUP_IS_NOT_HA
  The local consistency group does not have HyperSwap definitions.
- TARGET_NOT_CONNECTED
  There is currently no connection to the target system.
- REMOTE_TARGET_NOT_CONNECTED
  There is currently no connection from the target system.
- LOCAL_PEER_IS_NOT_MASTER
  The local peer is not primary.
- VOLUME_BELONGS_TO_HA_CONS_GROUP
  This volume belongs to a HyperSwap consistency group. The operation can be carried out on a volume which is not part of a consistency group, or on the consistency group itself.
- HA_IS_NOT_OPERATIONAL
  This HyperSwap relation is not operational. The operation cannot be carried out on a non-operational HyperSwap relation.
- HA_HIGH_AVAILABILITY_IS_ALREADY_DISABLED
  Automatic failover is already disabled (valid only for ha_high_availability_disable).

Converting a HyperSwap relation into a sync mirror

Use the **ha_convert_into_mirror** command to change a HyperSwap relation into a sync mirror.

```
ha_convert_into_mirror < vol=VolName | cg=cgName >
```

Parameters

Name	Type	Description	Mandatory
vol	Object name	Master volume.	N
cg	Object name	Master consistency group name or a list of master consistency groups.	N

Once the HyperSwap relation is converted into a sync mirror, it will no longer be retrieved by the ha_list command. Instead, it will appear in the output of the mirror_list command.

Prerequisites:

- The Slave volume is not mapped
- The local peer role is Master
- If carried out on a Slave volume, the HyperSwap connectivity must be down (verified with the target_list command)

The operation modifies the Slave volume SCSI identity. The change affects the WWN and the serial, while the volume name, external ID, and metadata remain unchanged.

Once the HyperSwap relation is converted into a sync mirror, it will no longer be retrieved by the ha_list command. Instead, it will appear in the output of the mirror_list command.

Example:

```
ha_convert_into_mirror vol=regular_volume
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.
- CONS_GROUP_BAD_NAME
  The consistency group name does not exist.
- VOLUME_IS_NOT_HA
  This is not a HyperSwap volume.
- CONS_GROUP_IS_NOT_HA
  The local consistency group does not have HyperSwap definitions.
- TARGET_NOT_CONNECTED
  There is currently no connection to the target system.
- REMOTE_TARGET_NOT_CONNECTED
  There is currently no connection from the target system.
- LOCAL_PEER_IS_NOT_MASTER
  The local peer is not primary.
- SLAVE_VOLUME_IS_MAPPED
  The secondary volume is mapped.
- VOLUME_BELONGS_TO_HA_CONS_GROUP
  This volume belongs to a HyperSwap consistency group. The operation can be carried out on a volume which is not part of a consistency group, or on the consistency group itself.
- HA_IS_NOT_OPERATIONAL
  This HyperSwap relation is not operational. The operation cannot be carried out on a non-operational HyperSwap relation.
- VOLUME_HAS_MULTIPLE_RELATIONS
  The volume has multiple relations. The operation is not allowed or a target must be specified.
- CONS_GROUP_HAS_MULTIPLE_RELATIONS
  The consistency group has multiple relations. The operation is not allowed or a target must be specified.

Converting a sync mirror into a HyperSwap relation

Use the **mirror_convert_into_ha** command to change a sync mirror into a HyperSwap relation.

```
mirror_convert_into_ha < vol=VolName | cg=cgName >
```

Parameters

Name	Type	Description	Mandatory
vol	Object name	Master volume.	N
cg	Object name	Master consistency group name or a list of master consistency groups.	N

Prerequisites

- The remote volume is not mapped
- The mirror is a sync mirror
- The local mirror is the Master
- The target connectivity is active and sufficient for HyperSwap (verified with the target_list command)
- The peers are connected to a properly configured and active Quorum Witness
- Both volumes have the same name

The command is issued on the Master peer and affects both peers, that are converted to HyperSwap at the same time. The HyperSwap functionality requires certain configuration elements, such as Quorum Witness, and will fail if they do not exist or are not in the required state.

Once the command has successfully completed, the relation will no longer be retrieved by the mirror_list command, but rather by the ha_list command.

The conversion process usually continues after the command completion event has been issued. While it continues, the indication of the high availability activation state Automatic Failover is Inactive, because metadata is still not synchronized. Once the conversion has completed, the HyperSwap volume can be mapped from the remote system and the new mapping can be used to perform I/O to that volume.

Example:

```
mirror_convert_into_ha vol=regular_volume
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.
- BAD_REMOTE_VOLUME_NAME
  The secondary volume name does not exist.
- CONS_GROUP_BAD_NAME
  The consistency group name does not exist.
- VOLUME_NO_MIRROR
  The local volume does not have remote mirroring definitions.
- CONS_GROUP_NO_MIRROR
  The local consistency group does not have remote mirroring definitions.
- MIRROR_IS_STANDBY
  The mirror is marked as Standby.
- TARGET_NOT_CONNECTED
  There is currently no connection to the target system.
- REMOTE_TARGET_NOT_CONNECTED
  There is currently no connection from the target system.
- LOCAL_PEER_IS_NOT_MASTER
  The local peer is not primary.

- SLAVE_VOLUME_IS_MAPPED
  The secondary volume is mapped.
- VOLUME_BELONGS_TO_MIRRORED_CONS_GROUP
  The volume mirror is part of a consistency group mirror.
- TARGET_HAS_NO_QUORUM_WITNESS
  The local target does not have a Quorum Witness defined.
- REMOTE_TARGET_HAS_NO_QUORUM_WITNESS
  The target on the remote system does not have a Quorum Witness defined.
- HA_LOCAL_PEER_HAS_NO_QUORUM_WITNESS_CONNECTIVITY
  The local peer connectivity to the Quorum Witness is not operational.
- HA_INCOMPATIBLE_TARGET_VERSION
  The automatic failover is not supported between the system versions of the specified peers.
- MIRROR_TYPE_IS_NOT_SYNC
  Mirror type is not Sync Best Effort.
- MIRROR_IS_NON_OPERATIONAL
  The mirror is non-operational.
- MIRROR_RETRY_OPERATION
  There is an operation in progress on this mirror.

  Troubleshooting: Retry the command in a few seconds.
- REMOTE_MAX_HA_REACHED
  The maximum number of HyperSwap relationships is already reached on the remote machine.
- LOCAL_MAX_HA_REACHED
  The maximum number of HyperSwap relationships is already reached on the local machine.
- VOLUME_HAS_OLVM
  An IBM Hyper-Scale Mobility relationship is defined for this volume.
- VOLUME_HAS_HA
  This operation is forbidden on a volume with a HyperSwap relation.
- MAX_NUM_OF_PROXY_VOLUME_REACHED
  Failed to create a mirror, because the maximum number of proxy volumes is exceeded.
- TARGET_VOLUME_HAS_OLVM
  This target volume is part of an IBM Hyper-Scale Mobility relationship.
- TARGET_VOLUME_HAS_HA
  This operation is forbidden, if the target volume is a peer in a HyperSwap relation.
- REMOTE_MAX_METADATA_OBJECTS_REACHED
  The maximum number of metadata objects has been reached on a remote system.
- LOCAL_AND_REMOTE_VOLUME_NAMES_ARE_DIFFERENT
  Local and remote volume names are different.
- HA_TARGET_QUORUM_WITNESS_IS_NOT_ACTIVATED
  The Quorum Witness associated with the target is deactivated.
- TARGET_PEER_NOT_HEALTHY
  The target peer is not identified as healthy.

  Troubleshooting: Check the Quorum Witness configuration.
- HA_CONNECTIVITY_NOT_SUFFICIENT
  The connectivity between the systems is not sufficient for the automatic failover.
- HA_PEER_QUORUM_WITNESS_CONFIGURATION_NOT_VERIFIED
  Connectivity between the peer and the Quorum Witness is not verified.
- HA_REMOTE_PEER_QUORUM_WITNESS_CONFIGURATION_NOT_VERIFIED
  Connectivity between the remote peer and the Quorum Witness is not verified.
- HOST_TYPE_IS_NOT_CONFIGURED
  Cannot associate a HyperSwap volume with a host of unconfigured type. IMPORTANT: Read the HyperSwap chapter in the 'Best Practices' document to understand the solution requirements.
- VOLUME_HAS_MULTIPLE_RELATIONS
  The volume has multiple relations. The operation is not allowed or a target must be specified.
- CONS_GROUP_HAS_MULTIPLE_RELATIONS
  The consistency group has multiple relations. The operation is not allowed or a target must be specified.
- LOCAL_VOLUME_HAS_TOO_MANY_METADATA_OBJECTS

Local volume has too many metadata objects.

Troubleshooting: Contact IBM Support.

Creating a new Quorum Witness

Use the **quorum_witness_define** command to create a new Quorum Witness definition in the system and connect the system to the Quorum Witness.

```
quorum_witness_define name=qw_name certificate=qw_certificate address=qw_address [ port=qw_port ]
 [ activate=<yes|no> ]
```

Parameters

Name	Type	Description	Mandatory	Default
name	Object name	The name of the Quorum Witness to be created.	Y	N/A
certificate	N/A	The public certificate or certificate chain of the Quorum Witness (see below for details).	Y	N/A
address	N/A	The Quorum Witness address: IPv4, IPv6 (full format only) or DNS name.	Y	N/A
port	Positive integer	The port used for Quorum Witness communications.	N	8460
activate	Boolean	Defines whether to activate the Quorum Witness upon creation.	N	yes

This command defines a Quorum Witness to be used for HyperSwap relations. Up to two Quorum Witness instances may be defined. Setting the activate parameter to yes will also activate the Quorum Witness and start the communication between the system and the Quorum Witness.

The certificate parameter:

The value of the certificate parameter is the content of a PEM file with asterisks instead of newlines. Chained certificates are supported. The total maximal length of a PEM file holding chained certificates (leaf first, root last) is 15360 characters (including the asterisk characters). In Windows, you can drag-and-drop a PEM file from the Windows Explorer to the appropriate location in the CLI session window; the content will be added automatically.

Example:

```
quorum_witness_define name=q1 certificate="Certificate" address=192.0.2.1
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- QUORUM_WITNESS_NAME_ALREADY_EXISTS
  The Quorum Witness name already exists.
- QUORUM_WITNESS_ADDRESS_ALREADY_EXISTS
  The Quorum Witness IPv4/6 address or DNS name with same port number already exists.

  Troubleshooting: Use the already configured Quorum Witness or change the address or port number.
- QUORUM_WITNESS_MISSING_SERVICE_CERTIFICATE
  No valid certificate is defined for the Quorum Witness service.

  Troubleshooting: Use PKI commands to define the certificate for the Quorum Witness service.
- MAX_QUORUM_WITNESSES_REACHED
  Too many Quorum Witness instances are defined.

  Troubleshooting: Delete one or more Quorum Witness instances and then try again.
- MAX_ACTIVE_QUORUM_WITNESSES_REACHED
  The maximum number of active Quorum Witness instances is already reached.

  Troubleshooting: Deactivate an active Quorum Witness and then try again.
- SSL_CERTIFICATE_HAS_EXPIRED
  The SSL certificate has expired.
- SSL_CERTIFICATE_ISSUER_NOT_FOUND

The SSL certificate issuer was not found in the certificate chain.

- SSL_CERTIFICATE_INVALID_FORMAT
  The SSL certificate format is invalid or corrupted.
- SSL_CERTIFICATE_CHAIN_EMPTY
  No certificates were found in the input.
- SSL_CERTIFICATE_VERIFICATION_INTERNAL_ERROR
  The SSL certificate verification has failed because of an internal system error.
- SSL_CERTIFICATE_VERIFICATION_FAILED
  The SSL certificate chain verification failed.
- SSL_CERTIFICATE_NOT_YET_VALID
  The SSL certificate is not yet valid.

Listing Quorum Witnesses

Use the **quorum_witness_list** command to list all Quorum Witnesses defined in the system, or only the specified one.

```
quorum_witness_list [ name=qw_name ] [ domain=DomainName ]
```

Parameters

Name	Type	Description	Mandatory	Default
name	Object name	The internal name of the Quorum witness.	N	All Quorum Witnesses
domain	Object name	The domain name.	N	All domains

The command output includes the following fields:

- name: The Quorum Witness internal name
- quorum_id: A globally unique Quorum Witness ID
- address: The v6 or DNS address for communicating with the Quorum Witness
- port: The port for communicating with the Quorum Witness
- state: The state of the Quorum Witness in the system. The following values are available: **Activating**, **Activated**, **Deactivating**, and **Deactivated**.
- connection: The state of the connection with the Quorum Witness, accumulated across all Quorum Node connection statuses.
- external_name: The Quorum Witness external name
- db_health: The state of the Quorum Witness DB health. The following values are available:

Value	Meaning
OK	N/A
Recovery needed	A problem was identified in the Quorum Witness DB, that prevents it from operating properly.
Restore pending	The Quorum Witness DB was initialized. Once the Quorum Witness is activated, the system will automatically re-register to the Quorum Witness, and restore the information pertaining to its HyperSwap relations.

The following optional fields can be listed by explicitly specifying the proper columns:

- heartbeats_ok: Indicates whether heartbeats between the system and the Quorum Witness are properly sent and received
- secure_connection: The state of the security established on the connection. The possible values are:

Value	Meaning
None	Security is disabled on the Quorum Witness connectivity.
Verified	The connection security is verified.
Unauthorized	The system certificate was rejected by the Quorum Witness.

- version: The Quorum Witness software version
- id
- certificate
- db_init: The timestamp (in mono-time) of the last DB initialization
- first_event_id: The ID of the first relevant event
- last_event_id: The ID of the last relevant event

Example:

```
quorum_witness_list name=q1
```

Output:

```
Name     ID                 Address     Port   State       Connection    External Name
------   ----------------   ----------  -----  ----------  ----------    --------------
q1       quorum_witness_id  192.0.2.1   4567   Activated   Up            quorum_witness

Cont:
Db Health
---------
OK
```

Field ID	Field output	Default position
name	Name	1
quorum_id	ID	2
address	Address	3
port	Port	4
state	State	5
connection	Connection	6
external_name	External Name	7
db_health	Db Health	8
heartbeats_ok	Heartbeating	N/A
use_secure	Use Secure	N/A
secure_connection	Secure Connection	N/A
version	Version	N/A
id	ID	N/A
certificate	Certificate	N/A
db_init	DB Init	N/A
first_event_id	First Event Id	N/A
last_event_id	Last Event Id	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Updating a Quorum Witness definition

Use the **quorum_witness_update** command to update a Quorum Witness definition.

```
quorum_witness_update name=qw_name [ certificate=new_qw_certificate ] [ address=new_qw_address ]
 [ port=new_qw_port ]
```

Parameters

Name	Type	Description	Mandatory	Default
name	Object name	The Quorum Witness internal name.	Y	N/A
certificate	N/A	The public certificate or certificate chain of the Quorum Witness (see below for details).	N	Current value.
address	N/A	The Quorum Witness address: IPv4, IPv6 (full format only) or DNS name.	N	Current value.
port	Positive integer	A new communication port of the Quorum Witness.	N	Current value.

As a prerequisite for successfully completing this command, the Quorum Witness must be deactivated.

The certificate parameter:

The value of the certificate parameter is the content of a PEM file with asterisks instead of newlines. Chained certificates are supported. The total maximal length of a PEM file holding chained certificates (leaf first, root last) is 15360 characters (including the asterisk characters). In Windows, you can drag-and-drop a PEM file from the Windows Explorer to the appropriate location in the CLI session window; the content will be added automatically.

Example:

```
quorum_witness_update name=q1 address=192.0.2.1
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- QUORUM_WITNESS_BAD_NAME
  The Quorum Witness name does not exist.
- QUORUM_WITNESS_ADDRESS_ALREADY_EXISTS
  The Quorum Witness IPv4/6 address or DNS name with same port number already exists.

  Troubleshooting: Use the already configured Quorum Witness or change the address or port number.
- CANT_UPDATE_ACTIVATED_QUORUM_WITNESS
  Cannot update an activated Quorum Witness.
- SSL_CERTIFICATE_HAS_EXPIRED
  The SSL certificate has expired.
- SSL_CERTIFICATE_ISSUER_NOT_FOUND
  The SSL certificate issuer was not found in the certificate chain.
- SSL_CERTIFICATE_INVALID_FORMAT
  The SSL certificate format is invalid or corrupted.
- SSL_CERTIFICATE_CHAIN_EMPTY
  No certificates were found in the input.
- SSL_CERTIFICATE_VERIFICATION_INTERNAL_ERROR
  The SSL certificate verification has failed because of an internal system error.
- SSL_CERTIFICATE_VERIFICATION_FAILED
  The SSL certificate chain verification failed.
- SSL_CERTIFICATE_NOT_YET_VALID
  The SSL certificate is not yet valid.

Renaming a Quorum Witness

Use the **quorum_witness_rename** command to rename a Quorum Witness.

```
quorum_witness_rename name=qw_name new_name=qw_new_name
```

Parameters

Name	Type	Description	Mandatory
name	Object name	The Quorum Witness internal name.	Y
new_name	Object name	A new internal name of the Quorum Witness.	Y

Example:

```
quorum_witness_rename name=q1 new_name=q2
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- QUORUM_WITNESS_BAD_NAME
  The Quorum Witness name does not exist.
- QUORUM_WITNESS_NAME_ALREADY_EXISTS
  The Quorum Witness name already exists.

Deleting a Quorum Witness

Use the **quorum_witness_delete** command to delete a Quorum Witness.

```
quorum_witness_delete name=qw_name
```

Parameters

Name	Type	Description	Mandatory
name	Object name	The Quorum Witness internal name.	Y

The command can be issued only for a Quorum Witness that is not in use (attached to any target) and is deactivated or is being deactivated, that is its state is *Deactivated* or *Deactivating*.

Example:

```
quorum_witness_delete name=q1 -y
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_DELETE_QUORUM_WITNESS
  Are you sure you want to delete Quorum Witness *Quorum Witness*?

Return codes

- QUORUM_WITNESS_BAD_NAME
  The Quorum Witness name does not exist.
- CANT_DELETE_AN_ACTIVATED_QUORUM_WITNESS
  Cannot delete an activated Quorum Witness.

  Troubleshooting: Deactivate the Quorum Witness and then try again.
- CANT_DELETE_QUORUM_WITNESS_IN_USE
  Cannot delete a Quorum Witness when it is in use by a target.

Activating a Quorum Witness

Use the **quorum_witness_activate** command to activate a Quorum Witness.

```
quorum_witness_activate name=qw_name
```

Parameters

Name	Type	Description	Mandatory
name	Object name	The Quorum Witness internal name.	Y

This commands activates a Quorum Witness and starts heartbeat and status communication between the system and the Quorum Witness. The command is asynchronous, its completion is tracked by return codes.

Example:

```
quorum_witness_activate name=q1
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed

User Category	Permission
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- MAX_ACTIVE_QUORUM_WITNESSES_REACHED
  The maximum number of active Quorum Witness instances is already reached.

  Troubleshooting: Deactivate an active Quorum Witness and then try again.
- QUORUM_WITNESS_BAD_NAME
  The Quorum Witness name does not exist.
- QUORUM_WITNESS_IS_ALREADY_ACTIVE
  The Quorum Witness is already active or is being activated.

  Troubleshooting: Wait for the activation process to complete. If this issue persists, contact IBM Support.
- QUORUM_WITNESS_MISSING_SERVICE_CERTIFICATE
  No valid certificate is defined for the Quorum Witness service.

  Troubleshooting: Use PKI commands to define the certificate for the Quorum Witness service.
- SSL_CERTIFICATE_HAS_EXPIRED
  The SSL certificate has expired.

Deactivating a Quorum Witness

Use the **quorum_witness_deactivate** command to deactivate a Quorum Witness.

```
quorum_witness_deactivate name=qw_name
```

Parameters

Name	Type	Description	Mandatory
name	Object name	The Quorum Witness internal name.	Y

This commands deactivates a Quorum Witness and stops the heartbeat and status communication between the system and the Quorum Witness. It is not allowed to deactivate a Quorum Witness that is in use by an active HyperSwap relation.

Example:

```
quorum_witness_deactivate name=q1 -y
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_DEACTIVATE_QUORUM_WITNESS
  Are you sure you want to deactivate Quorum Witness *Quorum Witness*?

Return codes

- QUORUM_WITNESS_BAD_NAME
  The Quorum Witness name does not exist.
- QUORUM_WITNESS_IS_ALREADY_INACTIVE

The Quorum Witness is already inactive or is being deactivated.

Troubleshooting: Wait for the deactivation process to complete. If this issue persists, contact IBM Support.

- CANT_DEACTIVATE_QUORUM_WITNESS_IN_USE
  Cannot deactivate a Quorum Witness that is in use by a target with an active HyperSwap relationship.

Listing the Quorum Witness connection status

Use the **quorum_witness_connections_list** command to list the status of the connections with the Quorum Witness.

```
quorum_witness_connections_list [ name=qw_name ] [ module=module_id ]
```

Parameters

Name	Type	Description	Mandatory	Default
name	Object name	The internal name of the Quorum Witness to be listed.	N	All Quorum Witness connections.
module	N/A	The ID of the module to be listed.	N	All modules to which Quorum Witness is connected.

The command output includes the following fields:

- module_id. This is the ID of the module on which the Quorum Witness is running.
- name. This is the Quorum Witness internal name.
- connection. This is state of the connection with the Quorum Witness. The possible values are:

Value	Meaning
`Band content encoding`	N/A
`Cannot connect to peer`	N/A
`Cannot resolve the proxy address`	N/A
`Cannot resolve the host address`	N/A
`Connected`	A Quorum Node is successfully connected to the Quorum Witness.
`Connection timeout`	The Quorum Witness may be down or the address/port pair is incorrect.
`Failed to chunk data`	N/A
`Failed to initialize`	The HTTP connection initialization failed.
`HTTP communication error`	N/A
`HTTP Post communication error`	N/A
`Malformed URL`	N/A
`Out of memory`	N/A
`Quorum node has failed`	N/A
`Secure connection issue`	See details in the secure_connection field.
`Too many redirects`	N/A
`Transient error`	N/A

- secure_connection: The state of the security established on the connection. The possible values are:

Value	Meaning
`None`	Security is disabled on the Quorum Witness connectivity.
`CA certificate is unusable`	The provided Quorum Witness certificate is incorrect.
`Could not load CACERT file, missing or wrong format`	N/A
`Could not load CRL file, missing or wrong format`	N/A
`Failed initializing secure communication`	N/A
`Failed to shut down the SSL/TLS connection`	N/A
`General SSL/TLS failure`	N/A
`Issuer check failed`	N/A
`Peer verification failed`	N/A
`Problem with local certificate`	N/A
`Secure cipher error`	N/A
`Unauthorized`	The system certificate was rejected by the Quorum Witness.
`Verified`	The connection security is verified.

- heartbeating: Indicates whether the module heartbeats are successfully processed by the Quorum Witness.

Example:

```
quorum_connection_list name=q1 module=1:Module:3
```

Output:

```
Module ID    Name    Connection   Secure Connection   Heartbeating
------------ ------- ------------ ------------------- --------------
1:Module:3   q1      Connected    Verified            yes
```

Field ID	Field output	Default position
module_id	Module ID	1
name	Name	2
connection	Connection	3
secure_connection	Secure Connection	4
heartbeats_ok	Heartbeating	5

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Allowed

Getting Quorum Witness information

Use the **quorum_witness_info_get** command to get Quorum Witness information from a Quorum Witness.

```
quorum_witness_info_get name=qw_name
```

Parameters

Name	Type	Description	Mandatory
name	Object name	The Quorum Witness internal name.	Y

The command output includes the following fields:

- internal_name: The Quorum Witness internal name
- quorum_id: A globally unique Quorum Witness ID
- name: The Quorum Witness external name
- version: The Quorum Witness software version

The following optional fields can be listed by explicitly specifying the proper columns:

- db_state: JSON string (up to 256 bytes), representing the DB disk usage
- host_info: The type and the version of the host's operating system, on which the Quorum Witness is running
- protocol: The protocol version
- network_load: JSON string (up to 128 bytes), representing the count of connections and requests
- cpu_load: JSON string (up to 128 bytes), representing the CPU consumption percentage
- last_event_id: The ID of the last relevant event
- db_init: The timestamp (in mono-time) of the last DB initialization

Example:

```
quorum_witness_info_get name=q1
```

Output:

```
Name   ID                                External Name   Version
------ --------------------------------- --------------- ---------
q1     7f2f35834fea48b6b335aa09fd9179ae  FVT-QW104       1.0.0
```

Field ID	Field output	Default position
internal_name	Name	1
quorum_id	ID	2
name	External Name	3
version	Version	4
db_state	DB state	N/A
host_info	Host	N/A
protocol	Protocol Version	N/A
network_load	Network Load	N/A
cpu_load	CPU Load	N/A
last_event_id	Last Event ID	N/A
db_init	DB Initialization Timestamp	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Return codes

- QUORUM_WITNESS_BAD_NAME
  The Quorum Witness name does not exist.
- QUORUM_WITNESS_RESPONSE_TIMEOUT
  No response from Quorum Witness *Quorum Witness* within the designated timeout period.

Data migration commands

This section describes the command-line interface (CLI) for data migration.

- **Activating data migration**
  Use the **dm_activate** command to activate the data migration process.
- **Deactivating data migration**
  Use the **dm_deactivate** command to deactivate the data migration process.
- **Defining data migration configuration**
  Use the **dm_define** command to define a data migration configuration.
- **Deleting a data migration process**
  Use the **dm_delete** command to delete the data migration process.
- **Listing data migration statuses**
  Use the **dm_list** command to list data migration configuration and status.
- **Testing the data migration definition**
  Use the **dm_test** command to test the data migration configuration.

Activating data migration

Use the **dm_activate** command to activate the data migration process.

```
dm_activate vol=VolName
```

Parameters

Name	Type	Description	Mandatory
vol	Object name	The destination volume for data migration activation.	Y

This command activates the data migration process. This is either an initial activation or an activation after deactivation.

Upon activation, the data migration is tested in the same way as when using dm_test (see Testing the data migration definition), and this command fails if the data migration test fails.

This command has no effect if the process is already active.

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.
- VOLUME_NO_DM
  The local volume does not have data migration definitions.

- TARGET_NOT_CONNECTED
  There is currently no connection to the target system.
- REMOTE_VOLUME_NO_LUN
  The remote volume's LUN is unavailable.
- REMOTE_VOLUME_NO_READ_ACCESS
  The remote volume cannot be read.
- REMOTE_VOLUME_NO_WRITE_ACCESS
  The remote volume is write-protected.
- BAD_REMOTE_VOLUME_SIZE
  The primary and secondary volumes contain a different number of blocks.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support

Deactivating data migration

Use the **dm_deactivate** command to deactivate the data migration process.

```
dm_deactivate vol=VolName
```

Parameters

Name	Type	Description	Mandatory
vol	Object name	The local volume on which the data migration process is to be deactivated.	Y

Hosts are not served while the data migration process is inactive.

This command has no effect if the data migration process is already inactive.

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_DEACTIVATE_DATA_MIGRATION
  Deactivation will stop all applications. After deactivation, data migration can be deleted.
- ARE_YOU_SURE_YOU_WANT_TO_DEACTIVATE_SOURCE_UPDATING_DATA_MIGRATION
  Deactivation may cause loss of access to hosts, and will stop all applications. After deactivation, data migration can be deleted.

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.
- VOLUME_NO_DM
  The local volume does not have data migration definitions.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support

Defining data migration configuration

Use the **dm_define** command to define a data migration configuration.

```
dm_define vol=VolName target=TargetName lun=SourceLUN source_updating=<yes|no> [ create_vol=<yes|no> ]
 [ pool=PoolName ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	Data migration destination volume on the local system.	Y	N/A
target	Object name	Remote system containing the source volume.	Y	N/A
lun	Integer	LUN of the source volume.	Y	N/A
source_updating	Boolean	Specifies whether to use source volume updating.	Y	N/A
create_vol	Boolean	A Boolean that determines whether to create a new volume or to use an existing one.	N	No
pool	Object name	Name of the storage pool to contain the volume. Used only when creating a volume. Mandatory when creating a volume.	N	N/A

This command defines a data migration relation between a local volume and a remote volume. According to this definition, the local volume should reflect the remote volume.

After this configuration has been defined, it can be tested using the dm_test command (see Testing the data migration definition) and then activated using the dm_activate command (see Activating data migration). After this activation, hosts can read and write to this volume, and these operations are reflected on the remote volume.

The remote volume may be inaccessible when the command is executed. In this case, the definition is only used when data migration is tested.

The local system acts as a host to the remote system. The remote system should be configured to make the remote volume accessible to the local system through the specified LUN.

If source updating is specified, each write to the local volume is reflected as a write to the remote volume. Otherwise, writes on the local volume are not reflected and the remote volume is not changed.

The local volume must be formatted.

If create_vol is set to *yes*, the volume is created. In this case the size of the newly created volume is identical to the size of the source volume. When creating a volume, a pool name must be specified. Creating a volume fails if there is no connectivity to the target since the volume's size is unknown.

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- VOLUME_SIZE_VERY_LARGE_ARE_YOU_SURE
  The volume size is very large. It may not be possible to mirror this volume to older versions of the storage system. Are you sure?

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.
- TARGET_BAD_NAME
  The target name does not exist.
- VOLUME_IS_SNAPSHOT
  The operation is not permitted on snapshots.
- VOLUME_HAS_MIRROR
  A mirror is defined for this volume.
- VOLUME_BELONGS_TO_CG
  The volume belongs to a consistency group.
- VOLUME_HAS_DATA_MIGRATION
  Data Migration is defined for this volume.
- VOLUME_HAS_SNAPSHOTS
  The volume has snapshots.
- VOLUME_NOT_FORMATTED
  The local volume is not formatted.
- VOLUME_EXISTS
  The volume name already exists.
- POOL_DOES_NOT_EXIST

The storage pool does not exist.

- VOLUME_BAD_PREFIX
  The volume name has a reserved prefix.
- NOT_ENOUGH_SPACE
  No space to allocate for the volume's current usage.
- MAX_VOLUMES_REACHED
  The maximum allowed number of volumes is already reached.
- REMOTE_VOLUME_NO_LUN
  The remote volume's LUN is unavailable.
- TARGET_NOT_CONNECTED
  There is currently no connection to the target system.
- VOLUME_CANNOT_HAVE_ZERO_SIZE
  The volume size cannot be zero.
- ILLEGAL_LUN
  The LUN is out of range.
- TARGET_IS_MIRRORING
  The target machine is defined for remote mirroring only.
- NO_ONLINE_MIGRATION_WITHOUT_SOURCE_UPDATING
  Data migration without automatic migration must be defined as a source update.
- MIGRATION_ALREADY_DEFINED_FOR_LUN
  Data migration is already defined from lun *LUN* of target '*Target*'.
- VOLUME_SIZE_ABOVE_LIMIT
  The specified volume size is above the limit.
- INVALID_SLICE_OFFSET
  Slice offset is illegal.
- ENCRYPTION_IN_PROGRESS
  The system is in the process of changing the encryption activation state.
- DOMAIN_MAX_VOLUMES_REACHED
  The domain exceeds the maximum allowed number of volumes.
- MAX_DMS_REACHED
  The maximum number of remote volumes (mirror/migration) is already reached.

  Troubleshooting: Delete unnecessary data migration objects.
- DOMAIN_MAX_DMS_REACHED
  The domain exceeds the maximum allowed number of data migrations.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support
- VOLUME_TOO_BIG
  No space to allocate to the volume.

Deleting a data migration process

Use the **dm_delete** command to delete the data migration process.

```
dm_delete vol=VolName [ force_delete=<yes|no> ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	The name of the volume whose data migration process is to be deleted.	Y	N/A
force_delete	Boolean	When set to *yes*, forces the deletion even if data migration is not complete.	N	no

This command deletes a data migration process and its configuration after it is completed and the target has a full copy of the data. After deleting a data migration process, the local volume must be explicitly deleted manually.

Using the force_delete parameter

The force_delete parameter enables deletion of a data migration process when only the source has a complete copy of data. This is mostly necessary if the data migration configuration is incorrect: a wrong source volume is selected or the defined volume (block) size is unsuitable.

When data migration is source updating and, therefore, the source has full data, the force_delete parameter is allowed as well.

It is currently impossible to delete a data migration process when neither the source nor the target have complete data.

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.
- VOLUME_NO_DM
  The local volume does not have data migration definitions.
- DM_IS_NOT_SYNCHRONIZED
  The data migration process has not been completed.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support
- DM_OPERATION_NOT_ALLOWED_NEITHER_SOURCE_NOR_TARGET_CONTAIN_FULL_DATA
  Data migration was set up without source update and synchronization has not completed yet. As a result, neither the source volume nor the target volume contains a full image of the data.

Listing data migration statuses

Use the **dm_list** command to list data migration configuration and status.

```
dm_list [ vol=VolName ] [ domain=DomainName ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	The name of the volume to be listed.	N	All data migration volumes.
domain	Object name	The domain name.	N	All Domains

This command lists all data migration configuration and statuses, including the following information:

- Volume name
- Target name
- LUN
- Volume size (GB)
- Migration completed (GB)
- Migration activation (active/inactive)
- Migration status (synchronized, unsynchronized)
- Migration remaining (GB)
- Migration remaining (%)
- Estimated time to completion

Field ID	Field output	Default position
local_volume_name	Local Volume	1
target_name	Remote System	2
remote_volume_lun	Remote LUN	3
active	Active	4
sync_state	Status	5
connected	Target Connected	6
size_to_synchronize	Size To Sync (MiB)	N/A
operational	Operational	N/A
sync_progress	Sync Progress (%)	N/A
start_migration_automatically	Start Data Migration Automatically	N/A
arch	Remote Arch	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Testing the data migration definition

Use the **dm_test** command to test the data migration configuration.

```
dm_test vol=VolName
```

Parameters

Name	Type	Description	Mandatory
vol	Object name	Destination volume for data migration testing.	Y

Command return codes indicate the types of test failures that may occur. Once a test is successful, then data migration can be activated.

If source updating is not defined for this data migration, the writing is not tested.

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.
- VOLUME_NO_DM
  The local volume does not have data migration definitions.
- TARGET_NOT_CONNECTED
  There is currently no connection to the target system.
- REMOTE_VOLUME_NO_LUN
  The remote volume's LUN is unavailable.
- REMOTE_VOLUME_NO_READ_ACCESS
  The remote volume cannot be read.
- REMOTE_VOLUME_NO_WRITE_ACCESS
  The remote volume is write-protected.
- BAD_REMOTE_VOLUME_SIZE
  The primary and secondary volumes contain a different number of blocks.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support

IBM Hyper-Scale Mobility commands

This section describes the command-line interface (CLI) for IBM Hyper-Scale Mobility.

- **Creating an IBM Hyper-Scale Mobility relation**
  Use the **olvm_create** command to define an IBM Hyper-Scale Mobility configuration.

- **Activating a volume migration**
  Use the **olvm_activate** command to activate an IBM Hyper-Scale Mobility migration for a defined an IBM Hyper-Scale Mobility relation.
- **Deactivating IBM Hyper-Scale Mobility migration**
  Use the **olvm_deactivate** command to deactivate IBM Hyper-Scale Mobility migration for a defined IBM Hyper-Scale Mobility relation.
- **Aborting a defined or activated IBM Hyper-Scale Mobility process**
  Use the **olvm_abort** command to abort a defined or activated IBM Hyper-Scale Mobility process.
- **Moving the IBM Hyper-Scale Mobility source volume to a Proxy state**
  Use the **olvm_proxy** command to move the IBM Hyper-Scale Mobility source volume to a Proxy state.
- **Deleting an IBM Hyper-Scale Mobility relation**
  Use the **olvm_delete** command to delete an IBM Hyper-Scale Mobility relation and attributes.
- **Listing the IBM Hyper-Scale Mobility status**
  Use the **olvm_list** command to list the IBM Hyper-Scale Mobility configuration and status.

Creating an IBM Hyper-Scale Mobility relation

Use the **olvm_create** command to define an IBM Hyper-Scale Mobility configuration.

```
olvm_create < vol=VolName remote_pool=RemotePoolName > target=TargetName
```

Parameters

Name	Type	Description	Mandatory
vol	Object name	An IBM Hyper-Scale Mobility volume on the local system.	Y
target	Object name	Remote system containing the destination volume.	Y
remote_pool	Object name	Name of the storage pool to contain the destination volume.	Y

This command creates an IBM Hyper-Scale Mobility relation through identifying the source volume and the destination system and storage pool.

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- VOLUME_SIZE_VERY_LARGE_ARE_YOU_SURE
  The volume size is very large. It may not be possible to mirror this volume to older versions of the storage system. Are you sure?

Return codes

- VOLUME_NOT_APPLICABLE_FOR_OLVM
  The volume is not applicable to IBM Hyper-Scale Mobility.
- VOLUME_BAD_NAME
  The volume name does not exist.
- VOLUME_IS_SNAPSHOT
  The operation is not permitted on snapshots.
- TARGET_BAD_NAME
  The target name does not exist.
- VOLUME_IS_MASTER
  This local volume is already defined as a primary volume.
- TARGET_BAD_TYPE
  The target machine is not an XIV machine.
- TARGET_NO_ACCESS
  No access permissions to the secondary machine.
- TARGET_NOT_CONNECTED
  There is currently no connection to the target system.
- REMOTE_TARGET_NOT_CONNECTED
  There is currently no connection from the target system.
- MAX_MIGRATIONS_REACHED

The maximum number of migrations is already reached.

- REMOTE_MAX_MIGRATIONS_REACHED
  The maximum number of migrations is already reached on the remote machine.
- REMOTE_POOL_DOES_NOT_EXIST
  The pool does not exist on the remote machine.
- BAD_REMOTE_VOLUME_SIZE
  The primary and secondary volumes contain a different number of blocks.
- NOT_ENOUGH_SPACE_ON_REMOTE_MACHINE
  Not enough free space to set the requested size of the secondary volume.
- REMOTE_VOLUME_EXISTS
  The secondary volume with the indicated name already exists. The name cannot be reused.
- REMOTE_VOLUME_IS_MASTER
  A volume on the remote machine is already defined as primary.
- REMOTE_VOLUME_IS_SLAVE
  A volume on the remote machine is already defined as secondary.
- REMOTE_MAX_VOLUMES_REACHED
  The maximum number of volumes on the remote machine is already reached.
- TIMEOUT
  A remote operation was not completed in time.
- VOLUME_BAD_PREFIX
  The volume name has a reserved prefix.
- REMOTE_VOLUME_HAS_DATA_MIGRATION
  Data migration is already defined for the secondary volume.
- VOLUME_HAS_OLVM
  An IBM Hyper-Scale Mobility relationship is defined for this volume.
- VOLUME_HAS_HA
  This operation is forbidden on a volume with a HyperSwap relation.
- TARGET_VOLUME_HAS_OLVM
  This target volume is part of an IBM Hyper-Scale Mobility relationship.
- TARGET_VOLUME_HAS_HA
  This operation is forbidden, if the target volume is a peer in a HyperSwap relation.
- REMOTE_VOLUME_LOCKED
  The secondary volume is locked.
- VOLUME_HAS_MIRRORING_SNAPSHOTS
  The volume has snapshots created by a previous mirroring process.
- REMOTE_MAX_MIRROR_CAPACITY_REACHED
  The maximum capacity for mirrored volumes is already reached on the remote machine.
- TARGET_DOES_NOT_ACCEPT_XIV_COMMANDS
  The target system does not accept XIV management commands.
- MAX_VOLUMES_REACHED
  The maximum allowed number of volumes is already reached.
- VOLUME_LOCKED
  The volume is locked.
- NO_ASYNC_IN_THIN_PROVISIONED_POOL
  A thin-provisioned pool cannot contain volumes with asynchronous mirroring.
- BAD_REMOTE_VOLUME_NAME
  The secondary volume name does not exist.
- REMOTE_VOLUME_SIZE_ABOVE_LIMIT
  The specified volume size is above the limit of the remote machine.
- MIRROR_RETRY_OPERATION
  There is an operation in progress on this mirror.

  Troubleshooting: Retry the command in a few seconds.
- ELECTRONIC_LICENSE_NOT_APPROVED
  Operation blocked until Electronic license approval

  Troubleshooting: Please retrieve Electronic license version and accept it

- VOLUME_NOT_FORMATTED
  The local volume is not formatted.
- MIRROR_TYPE_INCOMPATIBLE_WITH_TARGET
  A mirror of this type is not supported between the system versions of the specified peers.
- OLVM_INCOMPATIBLE_TARGET_VERSION
  IBM Hyper-Scale Mobility is not supported between the system versions of the specified peers.
- VOLUME_TOO_BIG
  No space to allocate to the volume.
- VOLUME_SIZE_ABOVE_LIMIT
  The specified volume size is above the limit.
- INVALID_SLICE_OFFSET
  Slice offset is illegal.
- VOLUME_IS_OLVM_PROXY
  The volume is in an IBM Hyper-Scale Mobility Proxy phase.
- VOLUME_IS_SLAVE
  The volume is defined as a secondary volume.
- REMOTE_VOLUME_IS_SNAPSHOT
  The secondary volume is a snapshot.
- VOLUME_EXISTS
  The volume name already exists.
- SLAVE_VOLUME_NOT_FORMATTED
  The secondary volume is not formatted.
- VOLUME_BELONGS_TO_CG
  The volume belongs to a consistency group.
- VOLUME_HAS_DATA_MIGRATION
  Data Migration is defined for this volume.
- MAX_MIRRORS_REACHED
  The maximum number of mirrors is already reached.
- VOLUME_CANNOT_HAVE_ZERO_SIZE
  The volume size cannot be zero.
- ASYNC_MIRROR_REMOTE_RPO_TOO_SHORT
  The specified remote RPO is too short.
- POOL_DOES_NOT_EXIST
  The storage pool does not exist.
- REMOTE_VOLUME_NOT_APPLICABLE_FOR_OLVM
  The remote volume is not applicable to IBM Hyper-Scale Mobility.
- REMOTE_SCHEDULE_DOES_NOT_EXIST
  The specified schedule does not exist on the remote machine.
- OLVM_DOES_NOT_SUPPORT_ISCSI_TARGET
  IBM Hyper-Scale Mobility does not support ISCSI targets.
- ASYNC_MIRROR_REMOTE_RPO_TOO_LONG
  The specified remote RPO is too long.
- OPERATION_DENIED_OBJECT_MANAGED
  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.
- ENCRYPTION_IN_PROGRESS
  The system is in the process of changing the encryption activation state.
- MAX_OLVM_REACHED
  The maximum allowed number of IBM Hyper-Scale Mobility relationships is already reached.
- DOMAIN_MAX_MIRRORS_REACHED
  The domain exceeds the maximum allowed number of mirrors.
- REMOTE_DOMAIN_MAX_MIGRATIONS_REACHED
  The maximum number of migrations is already reached in the remote machine domain.
- DOMAIN_HAS_NO_ACCESS_TO_TARGET
  The domain has no access to the target.
- REMOTE_DOMAIN_HAS_NO_ACCESS_TO_TARGET
  The secondary machine domain has no access to the target.

- DOMAIN_MAX_VOLUMES_REACHED
  The domain exceeds the maximum allowed number of volumes.
- REMOTE_DOMAIN_MAX_VOLUMES_REACHED
  The maximum number of volumes in the remote machine domain is already reached.
- MULTISITE_MAX_NUM_OF_MIRRORS_REACHED
  Failed to create the relation, the maximum allowed number of relations is already exceeded.
- REMOTE_VOLUME_TWO_SYNC_MIRRORS_NOT_ALLOWED
  Two synchronous mirrors were detected on the remote volume. This is not allowed.
- REMOTE_VOLUME_IS_MIRROR_MASTER
  The volume is primary in a mirror relationship, and cannot be secondary!
- REMOTE_VOLUME_MIRROR_LOOP_DETECTED
  A mirror loop was detected on the remote volume. This means that there is a mirror on the remote system, whose target is this system. Therefore, you cannot create a mirror with this target on this system.
- REMOTE_VOLUME_MASTER_ASYNC_MIRROR_DETECTED
  An asynchronous primary mirror was detected on the remote volume. The operation not allowed.
- MULTISITE_INCOMPATIBLE_TARGET_VERSION
  Multi-site is not supported between the system versions of the specified peers.
- REMOTE_VOLUME_HAS_MIRRORING_SNAPSHOTS
  The remote volume has snapshots created by a previous mirroring process.
- VOLUME_HAS_MULTIPLE_RELATIONS
  The volume has multiple relations. The operation is not allowed or a target must be specified.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support
- REMOTE_DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier of the remote system is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support
- SYSTEM_OUT_OF_PHYSICAL_SPACE
  The operation not allowed while the system is out of physical space.
- REMOTE_SYSTEM_OUT_OF_PHYSICAL_SPACE
  The operation not allowed while the remote system is out of physical space.
- TARGET_CONFIGURATION_AND_CODE_VERSION_DO_NOT_SUPPORT_OPERATION
  The target configuration and code version do not support this operation.

Activating a volume migration

Use the **olvm_activate** command to activate an IBM Hyper-Scale Mobility migration for a defined an IBM Hyper-Scale Mobility relation.

```
olvm_activate vol=VolName
```

Parameters

Name	Type	Description	Mandatory
vol	Object name	IBM Hyper-Scale Mobility source volume.	Y

This command is issued on the source.

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- VOLUME_BAD_NAME

The volume name does not exist.

- COMMAND_NOT_SUPPORTED_FOR_OLVM_SOURCE_IN_THIS_STATE
  The source is in an unsupported IBM Hyper-Scale Mobility state.
- DOMAIN_MAX_VOLUMES_REACHED
  The domain exceeds the maximum allowed number of volumes.
- MIRROR_CONFIGURATION_ERROR
  The mirror's local configuration does not match its remote configuration.
- MIRROR_CONS_GROUP_MEMBERSHIP_MISMATCH
  The mirrored consistency group contains different volumes on the primary and secondary machines. This problem occurs whenever the cg_add_vol or cg_remove_vol commands were previously issued, and the primary machine did not receive an acknowledgment from the secondary machine until the command timed out, or due to any other unexpected failure.
- MIRROR_RETRY_OPERATION
  There is an operation in progress on this mirror.

  Troubleshooting: Retry the command in a few seconds.
- COMMAND_NOT_SUPPORTED_FOR_OLVM_DESTINATION_IN_THIS_STATE
  The destination is in an unsupported IBM Hyper-Scale Mobility state.
- MIRROR_SIZE_MISMATCH
  The secondary and primary volume sizes are different.
- REMOTE_VOLUME_IS_MASTER
  A volume on the remote machine is already defined as primary.
- REMOTE_MAX_VOLUMES_REACHED
  The maximum number of volumes on the remote machine is already reached.
- VOLUME_NOT_DEFINED_FOR_OLVM
  The volume does not have IBM Hyper-Scale Mobility definitions.
- OLVM_ALREADY_ACTIVE
  The IBM Hyper-Scale Mobility relationship is already active.
- MAX_VOLUMES_REACHED
  The maximum allowed number of volumes is already reached.
- REMOTE_MIRROR_IS_STANDBY
  The remote mirror is marked as Standby.
- VOLUME_HAS_MULTIPLE_RELATIONS
  The volume has multiple relations. The operation is not allowed or a target must be specified.
- REMOTE_DOMAIN_MAX_VOLUMES_REACHED
  The maximum number of volumes in the remote machine domain is already reached.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support
- REMOTE_DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier of the remote system is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support
- REMOTE_SYSTEM_OUT_OF_PHYSICAL_SPACE
  The operation not allowed while the remote system is out of physical space.
- OLVM_INCOMPATIBLE_TARGET_VERSION
  IBM Hyper-Scale Mobility is not supported between the system versions of the specified peers.

Deactivating IBM Hyper-Scale Mobility migration

Use the **olvm_deactivate** command to deactivate IBM Hyper-Scale Mobility migration for a defined IBM Hyper-Scale Mobility relation.

```
olvm_deactivate vol=VolName
```

Parameters

Name	Type	Description	Mandatory
vol	Object name	The source volume.	Y

This command is issued on the source.

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_DEACTIVATE_OLVM
  Are you sure you want to deactivate IBM Hyper-Scale Mobility?

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.
- REMOTE_VOLUME_IS_MASTER
  A volume on the remote machine is already defined as primary.
- COMMAND_NOT_SUPPORTED_FOR_OLVM_SOURCE_IN_THIS_STATE
  The source is in an unsupported IBM Hyper-Scale Mobility state.
- VOLUME_NOT_DEFINED_FOR_OLVM
  The volume does not have IBM Hyper-Scale Mobility definitions.
- COMMAND_NOT_SUPPORTED_FOR_OLVM_DESTINATION_IN_THIS_STATE
  The destination is in an unsupported IBM Hyper-Scale Mobility state.
- OLVM_ALREADY_INACTIVE
  The IBM Hyper-Scale Mobility relationship is already inactive.
- REMOTE_MIRROR_IS_STANDBY
  The remote mirror is marked as Standby.
- VOLUME_HAS_MULTIPLE_RELATIONS
  The volume has multiple relations. The operation is not allowed or a target must be specified.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support

Aborting a defined or activated IBM Hyper-Scale Mobility process

Use the **olvm_abort** command to abort a defined or activated IBM Hyper-Scale Mobility process.

```
olvm_abort < vol=VolName [ force_abort=<yes|no> | force_abort_on_destination=<yes|no> ] >
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	The source volume.	Y	N/A
force_abort	Boolean	Determines whether to delete an IBM Hyper-Scale Mobility relation on the source.	N	No
force_abort_on_destination	Boolean	Determine whether to delete an IBM Hyper-Scale Mobility relation on the destination.	N	No

This command is issued on the source and has the option to abort the IBM Hyper-Scale Mobility process either from the source or from the destination.

Once issued, the source and destination are no longer part of an IBM Hyper-Scale Mobility relation. IBM Hyper-Scale Mobility attributes are deleted.

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed

User Category	Permission
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_ABORT_OLVM_RELATIONSHIP_IN_THIS_PHASE
  Are you sure you want to abort the IBM Hyper-Scale Mobility relationship?
- ARE_YOU_SURE_YOU_WANT_TO_FORCE_ABORT_OLVM_RELATIONSHIP_IN_THIS_PHASE
  Are you sure you want to force abort the IBM Hyper-Scale Mobility relationship?

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.
- VOLUME_NOT_DEFINED_FOR_OLVM
  The volume does not have IBM Hyper-Scale Mobility definitions.
- COMMAND_NOT_SUPPORTED_FOR_OLVM_SOURCE_IN_THIS_STATE
  The source is in an unsupported IBM Hyper-Scale Mobility state.
- COMMAND_NOT_SUPPORTED_FOR_OLVM_DESTINATION_IN_THIS_STATE
  The destination is in an unsupported IBM Hyper-Scale Mobility state.
- OLVM_IS_ACTIVE
  The IBM Hyper-Scale Mobility relationship is active.
- FORCE_ABORT_NOT_ALLOWED
  A forced IBM Hyper-Scale Mobility abort is not allowed.
- VOLUME_HAS_MULTIPLE_RELATIONS
  The volume has multiple relations. The operation is not allowed or a target must be specified.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support

Moving the IBM Hyper-Scale Mobility source volume to a Proxy state

Use the **olvm_proxy** command to move the IBM Hyper-Scale Mobility source volume to a Proxy state.

```
olvm_proxy vol=VolName
```

Parameters

Name	Type	Description	Mandatory
vol	Object name	The source volume.	Y

This command moves the IBM Hyper-Scale Mobility source volume to a Proxy state where the source acts as a proxy to the destination.

The source becomes a proxy and the destination becomes the data 'owner'. Host writes are no longer written to the source and the volume data on the source is freed. The source volume and snapshot data are deleted.

This command is issued on the source.

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_OLVM_PROXY
  Are you sure you want to move the volume *Volume*to the Proxy state? The source volume and all volume snapshots will be deleted.

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.
- COMMAND_NOT_SUPPORTED_FOR_OLVM_SOURCE_IN_THIS_STATE
  The source is in an unsupported IBM Hyper-Scale Mobility state.
- COMMAND_NOT_SUPPORTED_FOR_OLVM_DESTINATION_IN_THIS_STATE
  The destination is in an unsupported IBM Hyper-Scale Mobility state.
- VOLUME_NOT_DEFINED_FOR_OLVM
  The volume does not have IBM Hyper-Scale Mobility definitions.
- OLVM_LINK_IS_NOT_UP
  The IBM Hyper-Scale Mobility link is not up. The mapping list cannot be updated.
- OLVM_PROXY_MOVE_INITIATED
  IBM Hyper-Scale Mobility volume move to the Proxy state has started.
- HOST_BAD_NAME
  The host name does not exist.
- ISCSI_HOST_ILLEGAL_PORT_NAME
  The port name for iSCSI Host is illegal.

  Troubleshooting: Port names for iSCSI Hosts must contain only printable characters.
- MAX_PORTS_REACHED
  The maximum number of ports defined in the system is already reached.
- HOST_PORT_EXISTS
  A host with this port ID is already defined.
- REMOTE_MAX_VIRTUAL_HOSTS_REACHED
  The maximum number of defined remote virtual hosts is already reached.
- OLVM_RETRY_OPERATION
  There is an operation in progress on this OLVM.

  Troubleshooting: Retry the command in a few seconds.
- VOLUME_HAS_MULTIPLE_RELATIONS
  The volume has multiple relations. The operation is not allowed or a target must be specified.
- MAX_METADATA_OBJECTS_REACHED
  The maximum number of metadata objects has been reached.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support
- LOCAL_VOLUME_HAS_TOO_MANY_METADATA_OBJECTS
  Local volume has too many metadata objects.

  Troubleshooting: Contact IBM Support.

Deleting an IBM Hyper-Scale Mobility relation

Use the **olvm_delete** command to delete an IBM Hyper-Scale Mobility relation and attributes.

```
olvm_delete vol=VolName [ force_delete=<yes|no> ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	The volume for IBM Hyper-Scale Mobility abort.	Y	N/A
force_delete	Boolean	Determines whether to delete an IBM Hyper-Scale Mobility relation on the destination.	N	No

This command is issued on the source. If there is no communication to the destination, the command can force delete the IBM Hyper-Scale Mobility relation.

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed

User Category	Permission
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_DELETE_OLVM_RELATIONSHIP_IN_THIS_PHASE
  Are you sure you want to delete the IBM Hyper-Scale Mobility relationship?
- ARE_YOU_SURE_YOU_WANT_TO_FORCE_DELETE_OLVM_RELATIONSHIP_IN_THIS_PHASE
  Are you sure you want to force delete the IBM Hyper-Scale Mobility relationship?

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.
- VOLUME_NOT_DEFINED_FOR_OLVM
  The volume does not have IBM Hyper-Scale Mobility definitions.
- COMMAND_NOT_SUPPORTED_FOR_OLVM_SOURCE_IN_THIS_STATE
  The source is in an unsupported IBM Hyper-Scale Mobility state.
- COMMAND_NOT_SUPPORTED_FOR_OLVM_DESTINATION_IN_THIS_STATE
  The destination is in an unsupported IBM Hyper-Scale Mobility state.
- FORCE_DELETE_NOT_ALLOWED
  A forced deletion of the IBM Hyper-Scale Mobility relation is not allowed.
- VOLUME_IS_MAPPED
  The volume mapped to a host cannot be deleted.
- VOLUME_HAS_MULTIPLE_RELATIONS
  The volume has multiple relations. The operation is not allowed or a target must be specified.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support

Listing the IBM Hyper-Scale Mobility status

Use the **olvm_list** command to list the IBM Hyper-Scale Mobility configuration and status.

```
olvm_list [ vol=VolName ] [ domain=DomainName ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	The volume name to be listed.	N	Displays details for IBM Hyper-Scale Mobility relations in the local system.
domain	Object name	The domain name.	N	All Domains

This command is issued on the source. The output includes the following information:

- Volume name
- Role (Source, Destination)
- Remote System
- Active (Yes, No)
- Phase (Migration, Proxy-Ready, Proxy)
- State
- Link Up

Field ID	Field output	Description	Default position
name	Volume name	N/A	1
role	Role	N/A	2
target_name	Remote System	N/A	3
active	Active	N/A	4
phase	Phase	N/A	5
state	State	N/A	6
connected	Link Up	N/A	7
sync_progress	Sync Progress (%)	N/A	N/A

Field ID	Field output	Description	Default position
size_to_synchronize	Size To Sync (MiB)	N/A	N/A
estimated_sync_time	Est. remaining time (sec)	N/A	N/A
mirror_error	Mirror Error	No Error, Secondary pool exhausted, Configuration error or No thin provisioning resources	N/A
arch	Remote Arch	N/A	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Event handling commands

This section describes the command-line interface (CLI) for event handling, including listing events, filtering and sending notifications.

- **Generating a custom event**
  Use the **custom_event** command to generate a custom event.
- **Generating a CSS product event**
  Use the **css_product_event** command to generate a CSS (Cloud Storage Solutions) custom event.
- **Defining a new event notification destination**
  Use the **dest_define** command to define a new destination for event notifications.
- **Deleting a destination**
  Use the **dest_delete** command to delete an event notification destination.
- **Listing event notification destinations**
  Use the **dest_list** command to list event notification destinations.
- **Renaming a destination**
  Use the **dest_rename** command to rename an event notification destination.
- **Testing a destination**
  Use the **dest_test** command to send a test message to an event notification destination.
- **Updating an event notification destination**
  Use the **dest_update** command to update a destination.
- **Adding a destination to a destination group**
  Use the **destgroup_add_dest** command to add an event notification destination to a destination group.
- **Creating a destination group**
  Use the **destgroup_create** command to create an event notification destinations group.
- **Updating an event notification destination group**
  Use the **destgroup_update** command to update a destination group.
- **Deleting a destination group**
  Use the **destgroup_delete** command to delete an event notification destination group.
- **Listing destination groups**
  Use the **destgroup_list** command to list destination groups.
- **Removing a destination from a destination group**
  Use the **destgroup_remove_dest** command to remove an event notification destination from a destination group.
- **Renaming a destination group**
  Use the **destgroup_rename** command to rename an event notification destination group.
- **Clearing alerting events**
  Use the **event_clear** command to clear alerting events.
- **Listing events**
  Use the **event_list** command to list system events.
- **Listing uncleared alerting events**
  Use the **event_list_uncleared** command to list uncleared alerting events.
- **Setting the threshold for event notification**
  Use the **event_redefine_threshold** command to redefine the threshold of a parameterized event.
- **Listing thresholds**
  Use the **event_threshold_list** to list event thresholds.
- **Activating a rule**
  Use the **rule_activate** command to activate an event notification rule.
- **Creating event notification rules**
  Use the **rule_create** command to create an event notification rule.
- **Deactivating a rule**
  Use the **rule_deactivate** command to deactivate an event notification rule.
- **Deleting event notification rules**
  Use the **rule_delete** command to delete an event notification rule.
- **Listing event notification rules**
  Use the **rule_list** command to list event notification rules.

- **Renaming event notification rules**
  Use the **rule_rename** command to rename an event notification rule.
- **Updating an event notification rule**
  Use the **rule_update** command to update an event notification rule.
- **Defining an SMS gateway**
  Use the **smsgw_define** command to define an SMS gateway.
- **Deleting an SMS gateway**
  Use the **smsgw_delete** command to delete an SMS gateway.
- **Listing SMS gateways**
  Use the **smsgw_list** command to list SMS gateways.
- **Prioritizing SMS gateways**
  Use the **smsgw_prioritize** command to set the priorities of the SMS gateways for sending SMS messages.
- **Renaming an SMS gateway**
  Use the **smsgw_rename** command to rename an SMS gateway.
- **Updating an SMS gateway**
  Use the **smsgw_update** command to update an SMS gateway.
- **Defining a new SMTP gateway**
  Use the **smtpgw_define** command to define an SMTP gateway.
- **Deleting an SMTP gateway**
  Use the **smtpgw_delete** command to delete the specified SMTP gateway.
- **Listing SMTP gateways**
  Use the **smtpgw_list** command to list SMTP gateways.
- **Prioritizing SMTP gateways**
  Use the **smtpgw_prioritize** command to prirotize SMTP gateways.
- **Renaming an SMTP gateway**
  Use the **smtpgw_rename** command to rename an SMTP gateway.
- **Updating an SMTP gateway**
  Use the **smtpgw_update** command to update the configuration of an SMTP gateway.
- **Generating an XMPNS admin control event**
  Use the **xmpns_admin_config_set** command to generate an XMPNS_ADMIN_CONTROL event.

Generating a custom event

Use the **custom_event** command to generate a custom event.

```
custom_event description=Description [ severity=<INFORMATIONAL|WARNING|MINOR|MAJOR|CRITICAL> ]
 [ internal=<yes|no> ]
```

Parameters

Name	Type	Description	Mandatory	Default
description	String	Description of the event.	Y	N/A
severity	N/A	Severity of the event.	N	Informational
internal	Boolean	Defines whether this is an XIV internal custom event.	N	no

This command can be used to either generate an event from a user application or host side software, or to test the event notification procedures.

Example:

```
custom_event description="Test started"
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Generating a CSS product event

Use the **css_product_event** command to generate a CSS (Cloud Storage Solutions) custom event.

```
css_product_event product=Product version=Version server=Server platform=Platform action=Action properties=Properties [
severity=<INFORMATIONAL|WARNING|MINOR|MAJOR|CRITICAL> ]
```

Parameters

Name	Type	Description	Mandatory	Default
product	String	Product name.	Y	N/A
version	String	Version information.	Y	N/A
server	String	Server name.	Y	N/A
platform	String	Platform information.	Y	N/A
action	String	Action information.	Y	N/A
properties	String	Properties information.	Y	N/A
severity	N/A	Severity of the event.	N	Informational

This command can be used to either generate an event from a user application or host side software, or to test the event notification procedures.

Example:

```
css_product_event product=product_name version=version_info server=server_info platform=platform_name action=action_name
properties=properties_details
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Defining a new event notification destination

Use the **dest_define** command to define a new destination for event notifications.

```
dest_define dest=DestName type=<SNMP|EMAIL|SMS|HTTPS> < snmp_manager=SNMPManager | < uri=HTTPSaddress [ proxy=
ProxyAddress [ proxy_port=ProxyPortNum ] ]
 > | email_address=email | <area_code=AreaCode    number=PhoneNumber> |    user=UserName> [
smtpgws=<SMTPGW1[,SMTPGW2]...|ALL> | smsgws=<SMSGW1[,SMSGW2]...|ALL> ]
 [ heartbeat_test_hour=HH:MM [ heartbeat_test_days=Day ] ] [ domain=DomainList ]
```

Parameters

Name	Type	Description	Mandatory	Default
dest	Object name	Destination name.	Y	N/A
type	Enumeration	Destination type for event notifications: be email, SMS, HTTPS or SNMP.	Y	N/A
snmp_manager	N/A	IP address or DNS name of the SNMP manager.	N	N/A
uri	N/A	IP address or DNS name of the HTTPS server. If a port different from the default should be used, specify it here.	N	N/A
proxy	N/A	IP address or DNS name of the proxy server to send HTTPS over.	N	None
proxy_port	Integer	Proxy port number to send HTTPS through. The default is 1080.	N	None
email_address	N/A	Email address.	N	N/A
smtpgws	Object name	List of SMTP gateways to be used.	N	ALL (all gateways).
area_code	N/A	Area code of the cellular number for SMS notification. Use digits, '-' or '.'	N	N/A
number	N/A	Cellular number for SMS notification. Use digits, '-' or '.'	N	N/A
smsgws	Object name	SMS gateways to be used for this destination.	N	ALL (all gateways).
user	Object name	User name, where the user's email or phone are used.	N	N/A
heartbeat_test_hour	N/A	The hour for periodic heartbeat testing in the format HH:MM	N	No heartbeat
heartbeat_test_days	N/A	List of days for heartbeat testing: a comma-separated list of 3-letter day names (such as "mon", "mon,fri", etc.).	N	No heartbeat
domain	N/A	Attach the destination to the specified domains. To define more than one domain, separated them with a comma. To specify all existing domains, use "*".	N	none

This command defines a destination for event notifications. There are four types of destinations: email, SMS, HTTPS and SNMP.

Note:
Defining the HTTPS destination can only be done upon an RPQ request.

- *Email* destinations are used for sending notifications via email. When defining a new destination of type Email, either the email address of the recipient must be specified in *email_address* or the user name must be specified in *user* (in this case the email address of that user is used).
- *SMS* destinations are used for sending notifications via SMS to cellular phones. When defining a new destination of type SMS, either the cellular phone number of the destination must be specified in *number* or the user name must be specified in *user* (in this case the cellular phone number of that user is used). To allow correct formatting, the area code must be separated from the local number.
- *SNMP* destinations are used for sending notifications by SNMP traps to SNMP managers. When defining a new destination of type SNMP, the IP address of the SNMP manager must be specified.
- *HTTPS* destinations are used for sending notifications to HTTPS servers. When defining a new destination of type HTTPS, the IP address of the HTTPS server must be specified.

By default, when sending an email notification, all SMTP gateways specified in the smtpgw_prioritize command (see Prioritizing SMTP gateways) are used, according to the order specified in that command. It is possible to define that sending emails to a specific destination will use specific SMTP gateway or gateways. This is done by specifying the smtpgws parameter.

The same logic applies to sending SMS messages. By default, SMS gateways specified in the smsgw_prioritize command (see Prioritizing SMS gateways) are used, according to the order specified in this command. It is possible to define that messages to a specific SMS destination will be sent through a specific SMS gateway or gateways.

Example:

```
dest_define dest=adminemail type=EMAIL email_address=storageadmin@example.com
```

Output:

```
Command completed successfully.
```

Example:

```
dest_define dest=monitoringserver type=SNMP snmp_manager=192.0.2.111
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- DEST_MAX_REACHED
  The maximum allowed number of destinations is already reached.
- DEST_NAME_ALREADY_EXISTS
  The destination name already exists.
- DEST_NAME_IS_DESTGROUP_NAME
  The destination name already exists as a destination group name.
- EMAIL_NOT_ALLOWED_FOR_DEST_TYPE
  This type of destination cannot have an email address.
- GATEWAY_NAME_APPEARS_TWICE
  The gateway name appears twice on the list.
- GATEWAY_NAME_DOES_NOT_EXIST
  The gateway name does not exist.
- SMSGWS_NOT_ALLOWED_FOR_DEST_TYPE
  This type of destination cannot have SMS gateways.
- SMTPGWS_NOT_ALLOWED_FOR_DEST_TYPE
  This type of destination cannot have SMTP gateways.
- SNMP_MANAGER_MUST_BE_SPECIFIED_FOR_DEST_TYPE
  This type of destination must have an SNMP manager.
- SNMP_MANAGER_NOT_ALLOWED_FOR_DEST_TYPE
  This type of destination cannot have an SNMP manager.
- NO_SMS_GATEWAYS_ARE_DEFINED

An SMS destination cannot be defined if no SMS gateways are defined.

- HTTPS_ADDRESS_NOT_ALLOWED_FOR_DEST_TYPE
  This type of destination cannot have an HTTPS address.
- PROXY_ADDRESS_NOT_ALLOWED_FOR_DEST_TYPE
  This type of destination cannot have a proxy address.
- SNMP_DESTS_CANNOT_REFER_TO_USERS
  An SNMP destination cannot refer to a user.
- HTTPS_DESTS_CANNOT_REFER_TO_USERS
  An HTTPS destination cannot refer to a user.
- NO_SMTP_GATEWAYS_ARE_DEFINED
  An email destination cannot be defined if no SMTP gateways are defined.
- USER_EMAIL_ADDRESS_IS_NOT_DEFINED
  The user's email address is not defined.
- USER_PHONE_NUMBER_IS_NOT_DEFINED
  The user's phone number is not defined.
- USER_NAME_DOES_NOT_EXIST
  The user name does not exist.
- INTERNAL_DESTS_CANNOT_REFER_TO_USERS
  An internal destination cannot refer to a user.
- DAY_APPEARS_TWICE
  The day '*Day*' appears twice in the list.

  Troubleshooting: Make sure that each day appears in the list only once.
- HTTPS_ADDRESS_MUST_BE_SPECIFIED_FOR_DEST_TYPE
  This type of destination must have an HTTPS address.
- DEST_TYPE_NOT_SUPPORTED
  This destination type is not supported.

  Troubleshooting: Contact IBM Support
- USER_IS_NOT_IN_DESTINATION_DOMAINS
  The user must be included in destination domains.
- DOMAIN_DOESNT_EXIST
  The domain does not exist.

Deleting a destination

Use the **dest_delete** command to delete an event notification destination.

```
dest_delete dest=DestName
```

Parameters

Name	Type	Description	Mandatory
dest	Object name	Name of the destination to be deleted.	Y

Destinations that are part of a destination group or used by a rule cannot be deleted.

Destinations cannot be deleted while there are uncleared alerting events.

Example:

```
dest_delete dest=itmanager
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed

User Category	Permission
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_DELETE_DESTINATION
  Are you sure you want to delete destination *Destination*?

Return codes

- DEST_NAME_DOES_NOT_EXIST
  The destination name does not exist.
- CANNOT_CHANGE_EVENT_CONF_WITH_ALERTING_EVENTS
  Cannot change an event configuration while there are alerting events.

  Troubleshooting: Clear all alerting events before changing an event configuration.
- DEST_IS_PART_OF_DESTGROUP
  The destination cannot be deleted because it is part of a destination group.
- DEST_APPEARS_IN_RULE
  The destination cannot be deleted because it appears in a rule.

  Troubleshooting: To delete the destination, first delete the rule.

Listing event notification destinations

Use the **dest_list** command to list event notification destinations.

```
dest_list [ dest=DestName ] [ type=<SNMP|EMAIL|SMS|HTTPS> ] [ domain=DomainName ] [ internal=<yes|no> ]
```

Parameters

Name	Type	Description	Mandatory	Default
dest	Object name	Destinations to be listed.	N	All destinations.
type	Enumeration	Filter only destinations of the specified type.	N	All types.
internal	Enumeration	Filter destinations by their internal XIV attribute.	N	no
domain	Object name	The domain name.	N	All Domains

This command lists the configuration of all defined destinations, or of a specific destination.

Field ID	Field output	Default position
name	Name	1
type	Type	2
email_address	Email Address	3
area_code	Area Code	4
number	Phone Number	5
snmp_manager	SNMP Manager	6
uri	HTTPS Address	7
gateways	Gateways	N/A
user	User	8
heartbeat_test_days	Heartbeat Days	N/A
heartbeat_test_hour	Heartbeat Time	N/A
creator	Creator	N/A
proxy	proxy server address	N/A
proxy_port	proxy port number	N/A

Example:

```
dest_list
```

Output:

```
Name               Type  Email Address                 Phone Number  Gateways
storagemanager     EMAIL storageadmin@example.com                    all
monitoringserver   SNMP
```

Access control

User Category	Permission	Condition
Storage administrator	Conditionally Allowed	Allowed, unless the internal parameter is specified.
Storage integration administrator	Disallowed	N/A
Application administrator	Conditionally Allowed	Allowed, unless the internal parameter is specified.
Security administrator	Disallowed	N/A
Read-only users	Conditionally Allowed	Allowed, unless the internal parameter is specified.
Technicians	Allowed	N/A

Renaming a destination

Use the **dest_rename** command to rename an event notification destination.

```
dest_rename dest=DestName new_name=Name
```

Parameters

Name	Type	Description	Mandatory
dest	Object name	The destination to be renamed.	Y
new_name	Object name	New name of the destination.	Y

Example:

```
dest_rename dest=adminemail new_name=storagemanager
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- CANNOT_CHANGE_EVENT_CONF_WITH_ALERTING_EVENTS
  Cannot change an event configuration while there are alerting events.

  Troubleshooting: Clear all alerting events before changing an event configuration.
- DEST_NAME_DOES_NOT_EXIST
  The destination name does not exist.
- DEST_NAME_IS_DESTGROUP_NAME
  The destination name already exists as a destination group name.
- DEST_NAME_ALREADY_EXISTS
  The destination name already exists.

Testing a destination

Use the **dest_test** command to send a test message to an event notification destination.

```
dest_test dest=DestName management_ip=IPaddress [ smtpgw=SMTPGatewayName ] [ smsgw=SMSGatewayName ]
 [ internal=<yes|no> ]
```

Parameters

Name	Type	Description	Mandatory	Default
dest	Object name	Name of the destination to be tested.	Y	N/A
management_ip	N/A	Management IP used for sending the event notification.	Y	N/A
smtpgw	Object name	SMTP gateway to be tested.	N	Default system choice.
smsgw	Object name	SMS gateway to be tested.	N	Default system choice.
internal	Boolean	Must be specified for XIV-internal destinations.	N	no

This command tests a destination by sending a test message, SMS or SNMP trap. Note that a successful return code from this command does not ensure notification delivery.

Some problems with SNMP, email, and SMS delivery may fail to be detected.

For email messages, the SMTP gateway must be specified (the destination is only tested through that gateway). The same applies to the SMS gateway.

Access control

User Category	Permission	Condition
Storage administrator	Conditionally Allowed	Allowed, unless the internal parameter is specified.
Storage integration administrator	Disallowed	N/A
Application administrator	Conditionally Allowed	Allowed, unless the internal parameter is specified.
Security administrator	Disallowed	N/A
Read-only users	Disallowed	N/A
Technicians	Allowed	N/A

Return codes

- DEST_NAME_DOES_NOT_EXIST
  The destination name does not exist.
- DEST_TEST_NOT_PERFORMED_SYSTEM_BUSY
  The test of destination '*Destination Name*' was not performed because the system is busy.

  Troubleshooting: Retry in a few seconds.
- GATEWAY_NAME_DOES_NOT_EXIST
  The gateway name does not exist.
- SMSGWS_MUST_BE_SPECIFIED_FOR_DEST_TYPE
  This type of destination must have SMS gateways.
- SMSGWS_NOT_ALLOWED_FOR_DEST_TYPE
  This type of destination cannot have SMS gateways.
- SMTPGWS_MUST_BE_SPECIFIED_FOR_DEST_TYPE
  This type of destination must have SMTP gateways.
- SMTPGWS_NOT_ALLOWED_FOR_DEST_TYPE
  This type of destination cannot have SMTP gateways.
- DEST_TEST_FAILED
  The test of destination '*Destination Name*' failed.
- SYSTEM_HAS_NO_SUCH_EXTERNAL_IP
  The system has no such external IP address.
- MODULE_CANNOT_SEND_MESSAGES
  The selected module cannot send messages.

  Troubleshooting: Contact IBM Support
- ONLY_TECHNICIAN_CAN_REFER_TO_INTERNAL_EVENT_OBJECTS
  Only technicians are allowed to refer to internal event objects.

Updating an event notification destination

Use the **dest_update** command to update a destination.

```
dest_update dest=DestName [ snmp_manager=SNMPManager ] [ uri=HTTPSaddress ] [ proxy=ProxyAddress ]
 [ proxy_port=ProxyPortNum ] [ email_address=email ]  [ smtpgws=<SMTPGW1[,SMTPGW2]...|ALL> ]
 [ area_code=AreaCode ] [ number=PhoneNumber ]  [ smsgws=<SMSGW1[,SMSGW2]...|ALL> ]
 [ user=UserName ] [ heartbeat_test_hour=HH:MM ] [ heartbeat_test_days=Day ]
 [ domain=DomainList ]
```

Parameters

Name	Type	Description	Mandatory	Default
dest	Object name	Destination name.	Y	N/A
snmp_manager	N/A	IP address or DNS name of the SNMP manager.	N	Keep unchanged.
uri	N/A	IP address or DNS name of HTTPS server.	N	Keep unchanged.

Name	Type	Description	Mandatory	Default
proxy	N/A	IP address or DNS name of proxy server to send HTTPS over.	N	Keep unchanged.
proxy_port	Integer	Proxy port number to send HTTPS through (1080 by default).	N	Keep unchanged.
domain	N/A	Attach the destination to the specified domains. To define more than one domain, separated them with a comma. To specify all existing domains, use "*".	N	Keep unchanged
email_address	N/A	Email address.	N	Keep unchanged.
smtpgws	Object name	List of SMTP gateways to be used.	N	Keep unchanged.
area_code	N/A	Area code of the cellular number for SMS notification.	N	Keep unchanged.
number	N/A	Cellular number for SMS notification.	N	Keep unchanged.
smsgws	Object name	SMS gateways to be used.	N	Keep unchanged.
user	Object name	User name, where the user's email or phone are used.	N	Keep unchanged.
heartbeat_test_hour	N/A	The hour of periodic heartbeat testing	N	Keep unchanged.
heartbeat_test_days	N/A	List of days for heartbeat testing: a comma-separated list of 3-letter day names (such as "mon", "mon,fri", and so on).	N	Keep unchanged.

The parameters of this command are identical to the Defining a new event notification destination command, except that the destination type cannot be changed. All relevant fields must be specified (not only the ones that are being changed).

Example:

```
dest_update dest=storagemanager email_address=admin@example.com
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- DEST_NAME_DOES_NOT_EXIST
  The destination name does not exist.
- AREA_CODE_MUST_BE_SPECIFIED_FOR_DEST_TYPE
  This type of destination must have an area code.
- AREA_CODE_NOT_ALLOWED_FOR_DEST_TYPE
  This type of destination cannot have an area code.
- CANNOT_CHANGE_EVENT_CONF_WITH_ALERTING_EVENTS
  Cannot change an event configuration while there are alerting events.

  Troubleshooting: Clear all alerting events before changing an event configuration.
- EMAIL_MUST_BE_SPECIFIED_FOR_DEST_TYPE
  This type of destination must have an email address.
- EMAIL_NOT_ALLOWED_FOR_DEST_TYPE
  This type of destination cannot have an email address.
- GATEWAY_NAME_APPEARS_TWICE
  The gateway name appears twice on the list.
- GATEWAY_NAME_DOES_NOT_EXIST
  The gateway name does not exist.
- NUMBER_MUST_BE_SPECIFIED_FOR_DEST_TYPE
  This type of destination must have a number.
- NUMBER_NOT_ALLOWED_FOR_DEST_TYPE

This type of destination cannot have a number.

- SMSGWS_NOT_ALLOWED_FOR_DEST_TYPE
  This type of destination cannot have SMS gateways.
- SNMP_MANAGER_NOT_ALLOWED_FOR_DEST_TYPE
  This type of destination cannot have an SNMP manager.
- NO_SMTP_GATEWAYS_ARE_DEFINED
  An email destination cannot be defined if no SMTP gateways are defined.
- DEST_CANNOT_HAVE_A_USER_AND_AN_EMAIL_ADDRESS
  The destination cannot simultaneously have an email address and refer to a user.
- DEST_CANNOT_HAVE_A_USER_AND_A_PHONE_NUMBER
  The destination cannot simultaneously have a phone number and refer to a user.
- USER_PHONE_NUMBER_IS_NOT_DEFINED
  The user's phone number is not defined.
- USER_NAME_DOES_NOT_EXIST
  The user name does not exist.
- INTERNAL_DESTS_CANNOT_REFER_TO_USERS
  An internal destination cannot refer to a user.
- DEST_HEARTBEAT_DAYS_BUT_NO_HOUR
  Destination heartbeat days are specified, but not heartbeat hour.
- HTTPS_ADDRESS_NOT_ALLOWED_FOR_DEST_TYPE
  This type of destination cannot have an HTTPS address.
- PROXY_ADDRESS_NOT_ALLOWED_FOR_DEST_TYPE
  This type of destination cannot have a proxy address.
- SNMP_DESTS_CANNOT_REFER_TO_USERS
  An SNMP destination cannot refer to a user.
- HTTPS_DESTS_CANNOT_REFER_TO_USERS
  An HTTPS destination cannot refer to a user.
- USER_EMAIL_ADDRESS_IS_NOT_DEFINED
  The user's email address is not defined.
- SMTPGWS_NOT_ALLOWED_FOR_DEST_TYPE
  This type of destination cannot have SMTP gateways.
- DAY_APPEARS_TWICE
  The day '*Day*' appears twice in the list.

  Troubleshooting: Make sure that each day appears in the list only once.
- SNMP_MANAGER_MUST_BE_SPECIFIED_FOR_DEST_TYPE
  This type of destination must have an SNMP manager.
- NO_SMS_GATEWAYS_ARE_DEFINED
  An SMS destination cannot be defined if no SMS gateways are defined.
- HTTPS_ADDRESS_MUST_BE_SPECIFIED_FOR_DEST_TYPE
  This type of destination must have an HTTPS address.
- DEST_TYPE_NOT_SUPPORTED
  This destination type is not supported.

  Troubleshooting: Contact IBM Support
- DOMAIN_DOESNT_EXIST
  The domain does not exist.
- USER_IS_NOT_IN_DESTINATION_DOMAINS
  The user must be included in destination domains.
- DESTINATION_IS_NOT_IN_RULE_DOMAINS
  The destination must be included in rule domains.
- DESTINATION_IS_NOT_IN_DESTGROUP_DOMAINS
  The destination must be included in the destination group domains.

Adding a destination to a destination group

Use the **destgroup_add_dest** command to add an event notification destination to a destination group.

```
destgroup_add_dest destgroup=GroupName dest=DestName
```

Parameters

Name	Type	Description	Mandatory
destgroup	Object name	Destination group name to which to add the destination.	Y
dest	Object name	Destination to be added to the group.	Y

The command fails if the destination group already contains the destination.

The command cannot be executed while there are uncleared alerting events.

Example:

```
destgroup_add_dest destgroup=alladmins dest=john
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- DESTGROUP_NAME_DOES_NOT_EXIST
  The destination group name does not exist.
- CANNOT_CHANGE_EVENT_CONF_WITH_ALERTING_EVENTS
  Cannot change an event configuration while there are alerting events.

  Troubleshooting: Clear all alerting events before changing an event configuration.
- DEST_NAME_DOES_NOT_EXIST
  The destination name does not exist.
- DESTGROUP_MAX_DESTS_REACHED
  The maximum allowed number of destinations is already reached in destination groups.
- DESTGROUP_ALREADY_INCLUDES_DEST
  The destination group already includes this destination name.
- DESTINATION_IS_NOT_IN_DESTGROUP_DOMAINS
  The destination must be included in the destination group domains.

Creating a destination group

Use the **destgroup_create** command to create an event notification destinations group.

```
destgroup_create destgroup=GroupName [ domain=DomainList ]
```

Parameters

Name	Type	Description	Mandatory	Default
destgroup	Object name	Destination group name.	Y	N/A
domain	N/A	Attach the destination group to the specified domains. To define more than one domain, separated them with a comma. To specify all existing domains, use "*".	N	none

This command creates a destination group, which is used by rules to send notifications to the entire group without specifying all the destinations for each rule. You can also add or remove destinations from the group, which eliminates the need to change the configuration of each rule separately.

Upon creation, the destination group is empty. To add a destination to a destination group, use the Adding a destination to a destination group command.

Example:

```
destgroup_create destgroup=alladmins
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- DESTGROUP_MAX_REACHED
  The maximum allowed number of destination groups is already reached.
- DESTGROUP_NAME_ALREADY_EXISTS
  The destination group name already exists.
- DESTGROUP_NAME_IS_DEST_NAME
  The destination group name already exists as a destination name.
- DOMAIN_DOESNT_EXIST
  The domain does not exist.

Updating an event notification destination group

Use the **destgroup_update** command to update a destination group.

```
destgroup_update destgroup=GroupName domain=DomainList
```

Parameters

Name	Type	Description	Mandatory
destgroup	Object name	Destination group name.	Y
domain	N/A	Attach the destination group to the specified domains. To define more than one domain, separated them with a comma. To specify all existing domains, use "*".	Y

Example:

```
destgroup_update destgroup=alladmins domain=D1,D2
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- DESTINATION_IS_NOT_IN_DESTGROUP_DOMAINS
  The destination must be included in the destination group domains.
- DOMAIN_DOESNT_EXIST
  The domain does not exist.
- DESTGROUP_IS_NOT_IN_RULE_DOMAINS
  The destination groups must be included in rule domains.

- DESTGROUP_NAME_DOES_NOT_EXIST
  The destination group name does not exist.
- INTERNAL_EVENT_OBJECTS_CANNOT_USE_SPECIFIC_DOMAINS
  Internal event objects cannot be defined on specific domains.

Deleting a destination group

Use the **destgroup_delete** command to delete an event notification destination group.

```
destgroup_delete destgroup=GroupName
```

Parameters

Name	Type	Description	Mandatory
destgroup	Object name	The name of the destination group to be deleted.	Y

Example:

```
destgroup_delete destgroup=dest1
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_DELETE_DESTINATION_GROUP
  Are you sure you want to delete destination group *Destination Group*?

Return codes

- CANNOT_CHANGE_EVENT_CONF_WITH_ALERTING_EVENTS
  Cannot change an event configuration while there are alerting events.

  Troubleshooting: Clear all alerting events before changing an event configuration.
- DESTGROUP_APPEARS_IN_RULE
  The destination group appears in a rule.

  Troubleshooting: To delete the destination group, first delete the rule.
- DESTGROUP_NAME_DOES_NOT_EXIST
  The destination group name does not exist.

Listing destination groups

Use the **destgroup_list** command to list destination groups.

```
destgroup_list [ destgroup=GroupName ] [ domain=DomainName ]
```

Parameters

Name	Type	Description	Mandatory	Default
destgroup	Object name	Destination group to be listed.	N	All groups.
domain	Object name	The domain name.	N	All Domains

This command lists all destination groups or a specific one. All the destinations are listed for each destination group.

Field ID	Field output	Default position
name	Name	1

Field ID	Field output	Default position
dests	Destinations	2
creator	Creator	N/A

Example:

```
destgroup_list
```

Output:

```
Name      Destinations
itstaff   john,michael,linda,monitoringserver
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Removing a destination from a destination group

Use the **destgroup_remove_dest** command to remove an event notification destination from a destination group.

```
destgroup_remove_dest destgroup=GroupName dest=DestName
```

Parameters

Name	Type	Description	Mandatory
destgroup	Object name	Group name.	Y
dest	Object name	Destination to be removed from the group.	Y

This command cannot be executed while there are uncleared alerting events.

Example:

```
destgroup_remove_dest destgroup=alladmins dest=john
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- DESTGROUP_NAME_DOES_NOT_EXIST
  The destination group name does not exist.
- DEST_NAME_DOES_NOT_EXIST
  The destination name does not exist.
- CANNOT_CHANGE_EVENT_CONF_WITH_ALERTING_EVENTS
  Cannot change an event configuration while there are alerting events.

  Troubleshooting: Clear all alerting events before changing an event configuration.
- DESTGROUP_DOES_NOT_INCLUDE_DEST
  The destination group does not include this destination name.

Renaming a destination group

Use the **destgroup_rename** command to rename an event notification destination group.

```
destgroup_rename destgroup=GroupName new_name=Name
```

Parameters

Name	Type	Description	Mandatory
destgroup	Object name	Destination group to be renamed.	Y
new_name	Object name	New name of the destination group.	Y

This command cannot be executed while there are uncleared alerting events.

Example:

```
destgroup_rename destgroup=alladmins new_name=itstaff
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- DESTGROUP_NAME_DOES_NOT_EXIST
  The destination group name does not exist.
- DESTGROUP_NAME_ALREADY_EXISTS
  The destination group name already exists.
- CANNOT_CHANGE_EVENT_CONF_WITH_ALERTING_EVENTS
  Cannot change an event configuration while there are alerting events.

  Troubleshooting: Clear all alerting events before changing an event configuration.
- DESTGROUP_NAME_IS_DEST_NAME
  The destination group name already exists as a destination name.

Clearing alerting events

Use the **event_clear** command to clear alerting events.

```
event_clear event_id=EventId [ all_preceding=<yes|no> ] [ internal=<yes|no|all> ]
```

Parameters

Name	Type	Description	Mandatory	Default
event_id	Positive integer	The ID number of the event to be cleared.	Y	N/A
all_preceding	Boolean	Clears all events preceding the specified event.	N	no
internal	Boolean	Clears XIV-internal events.	N	no

In order to ensure that an event was indeed received, an event notification may be sent repeatedly until it is cleared with a CLI command or from the GUI. Such events are called *alerting* events. An event is defined as *alerting* if at the time of the event's generation it was matched by an *alerting* rule, meaning a rule that has either snooze or escalation definitions.

Notifications for the alerting event are sent until it is cleared by this command. The clearing operation does not imply that the problem has been solved. It only implies that the event has been noted by the relevant person who takes responsibility for fixing the problem.

The user may clear either a specific event or all alerting events.

Example:

```
event_clear event_id=87
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission	Condition
Storage administrator	Conditionally Allowed	Allowed, unless the internal parameter is specified.
Storage integration administrator	Disallowed	N/A
Application administrator	Conditionally Allowed	Allowed, unless the internal parameter is specified.
Security administrator	Disallowed	N/A
Read-only users	Disallowed	N/A
Technicians	Allowed	N/A

Return codes

- ONLY_TECHNICIAN_CAN_REFER_TO_INTERNAL_EVENT_OBJECTS
  Only technicians are allowed to refer to internal event objects.

Listing events

Use the **event_list** command to list system events.

```
event_list [ max_events=MaxEventsToList ] [ after=TimeStamp ] [ before=TimeStamp ]
 [ min_severity=<INFORMATIONAL|WARNING|MINOR|MAJOR|CRITICAL> ] [ alerting=<yes|no|all> ]
 [ cleared=<yes|no|all> ] [ code=EventCode ]
 [ object_type=<cons_group|destgroup|dest|dm|host|map|mirror|pool|rule|smsgw|smtpgw|target|v
olume|cluster|ip_interface|ldap_conf|meta_data_object|sync_schedule|user|user_group|ldap_ser
ver|modules_status|elicense|ipsec_connection|ipsec_tunnel|cross_cons_group,...> ]
 [ internal=<yes|no|all> ] [ beg=BeginIndex ] [ end=EndIndex ] [ count_all=<yes|no> ]
 [ domain=DomainName ]
```

Parameters

Name	Type	Description	Mandatory	Default
max_events	Positive integer	Maximal number of events to list (see below for more details).	N	300
after	N/A	Earliest time/date.	N	No filter.
before	N/A	Latest time/date.	N	No filter.
min_severity	Enumeration	Minimum severity.	N	No filter.
alerting	Boolean	Filter alerting events.	N	No filter.
cleared	Boolean	Filter cleared events.	N	No filter.
code	N/A	Filter by a specific event code.	N	No filter.
object_type	Enumeration	Filter events by the type of the related system object.	N	No filter.
internal	Boolean	Filter XIV internal events.	N	No filter.
beg	Integer	Index of the first event to list. If negative, then counts from the end.	N	1
end	Integer	Index of the last event to list (not inclusive). If negative, then counts from the end.	N	Last event + 1.
count_all	Boolean	If yes, it scans all the events between beginning and end for computing the number of events meeting the criteria.	N	No.
domain	Object name	The domain name.	N	All Domains.

This command lists system events according to specified criteria, such as minimum severity, event type, and so on. The event list displays the following information for each event: timestamp, severity, code, user and description.

The max_events parameter

Events are listed and sorted by the time of creation in ascending order.

If the maximal number of events to display is limited by specifying the max_events parameter, and the number of retrieved events exceeds the maximum, only the most recent events will be displayed.

If the max_events parameter is not specified, the number of displayed events defaults to 300.

The fields **before** and **after**

The syntax for the **before** and **after** fields is as follows: **Y-M-D[.[h[:m[:s]]]]**, where the ranges are as follows:

- Y - year (four digit)
- M - month (1-12)
- D - day (1-31)
- h - hour (0-23, with 0 as default)
- m - minute (0-59, with 0 as default)
- s - second (0-59, with 0 as default)

The year, month and day are separated by dashes, while the optional hour, minute and second are separated by colons.

Example:

```
event_list max_events=10
```

Output:

```
Timestamp              Severity        Code
2009-05-12 15:10:16    Informational   START_WORK
2009-05-12 15:16:11    Informational   POOL_CREATE
2009-05-12 15:16:23    Informational   VOLUME_CREATE

Additional output fields
(lines are broken to fit the page width of this Guide):

User               Description
                   System has entered ON state.
xiv_development    Storage Pool of size 171GB was created with name
                   'p1_m'.
xiv_development    Volume was created with name 'master' and size 17GB in
                   Storage Pool with name 'p1_m'.
```

Field ID	Field output	Default position
timestamp	Timestamp	1
severity	Severity	2
code	Code	3
user_name	User	4
description	Description	5
index	Index	N/A
alerting	Alerting	N/A
cleared	Cleared	N/A
tshooting	Trouble Shooting	N/A

Access control

User Category	Permission	Condition
Storage administrator	Conditionally Allowed	Allowed, unless the internal parameter is specified.
Storage integration administrator	Conditionally Allowed	Allowed, unless the internal parameter is specified.
Application administrator	Conditionally Allowed	Allowed, unless the internal parameter is specified.
Security administrator	Conditionally Allowed	Allowed, unless the internal parameter is specified.
Read-only users	Conditionally Allowed	Allowed, unless the internal parameter is specified.
Technicians	Allowed	N/A

Return codes

- UNRECOGNIZED_EVENT_CODE
  '*String*' is not a recognized return code.

  Troubleshooting: Consult the manual for the list of valid return codes.
- CANNOT_READ_EVENTS
  Cannot read events.

  Troubleshooting: Contact IBM Support.
- DOMAIN_DOESNT_EXIST
  The domain does not exist.

Listing uncleared alerting events

Use the **event_list_uncleared** command to list uncleared alerting events.

```
event_list_uncleared [ domain=DomainName ]
```

Parameters

Name	Type	Description	Mandatory	Default
domain	Object name	The domain name.	N	All Domains

Example:

```
event_list_uncleared
```

Output:

```
Index   Code              Severity
```

```
------- --------------- ---------------
318     VOLUME_CREATE   Informational
666     VOLUME_DELETE   Informational
```

Field ID	Field output	Default position
index	Index	1
code	Code	2
severity	Severity	3

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Allowed
Read-only users	Allowed
Technicians	Allowed

Setting the threshold for event notification

Use the **event_redefine_threshold** command to redefine the threshold of a parameterized event.

```
event_redefine_threshold code=EventCode severity=<INFORMATIONAL|WARNING|MINOR|MAJOR|CRITICAL|NONE> threshold=<
ThresholdValue|NONE>
```

Parameters

Name	Type	Description	Mandatory
code	N/A	Event code.	Y
severity	Enumeration	Severity.	Y
threshold	Integer	Threshold value, or NONE to indicate that an event with this severity is not created.	Y

This command can be applied to parameterized events, that is events that are triggered when a certain parameter crosses a certain threshold. Using this command the user can change the threshold for event notification. Furthermore, multiple thresholds can be defined using multiple invocations of this command, one for each event severity. When the relevant parameter crosses a threshold, an event with the matching severity is created.

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- EVENT_DOES_NOT_HAVE_THRESHOLDS
  The event does not have thresholds.
- EVENT_THRESHOLD_IS_ILLEGAL
  An illegal value for the event threshold.

  Troubleshooting: Event threshold values must be monotonic.
- UNRECOGNIZED_EVENT_CODE
  '*String*' is not a recognized return code.

  Troubleshooting: Consult the manual for the list of valid return codes.
- LAST_EVENT_THRESHOLD_CANNOT_BE_DELETED
  The event must have at least one threshold value.

Listing thresholds

Use the **event_threshold_list** to list event thresholds.

```
event_threshold_list [ code=EventCode ]
```

Parameters

Name	Type	Description	Mandatory	Default
code	Enumeration	Filter by a specific event code.	N	no filter

Field ID	Field output	Default position
code	Code	1
has_thresholds	Has Thresholds?	N/A
not_in_use	Not In Use	N/A
replaced_by	Replaced By	N/A
default_thresholds.0	INFORMATIONAL(def)	7
default_thresholds.1	WARNING(def)	8
default_thresholds.2	MINOR(def)	9
default_thresholds.3	MAJOR(def)	10
default_thresholds.4	CRITICAL(def)	11
thresholds.0	INFORMATIONAL	2
thresholds.1	WARNING	3
thresholds.2	MINOR	4
thresholds.3	MAJOR	5
thresholds.4	CRITICAL	6

Example:

```
event_threshold_list
```

Output:

```
Code                                    INFORMATIONAL   WARNING   MINOR
--------------------------------------- --------------- --------- -------
STORAGE_POOL_SNAPSHOT_USAGE_INCREASED   none            80        90
STORAGE_POOL_VOLUME_USAGE_INCREASED     none            80        90

MAJOR   CRITICAL   INFORMATIONAL(def)   WARNING(def)   MINOR(def)   MAJOR(def)
------- ---------- -------------------- -------------- ------------ ------------
95      none       none                 80             90           95
95      none       none                 80             90           95

CRITICAL(def)
---------------
none
none
```

Access control

User Category	Permission	Condition
Storage administrator	Conditionally Allowed	Allowed, unless the internal parameter is specified.
Storage integration administrator	Disallowed	N/A
Application administrator	Conditionally Allowed	Allowed, unless the internal parameter is specified.
Security administrator	Disallowed	N/A
Read-only users	Conditionally Allowed	Allowed, unless the internal parameter is specified.
Technicians	Allowed	N/A

Activating a rule

Use the **rule_activate** command to activate an event notification rule.

```
rule_activate rule=RuleName
```

Parameters

Name	Type	Description	Mandatory
rule	Object name	The name of the rule to be activated.	Y

This command activates the specified rule. An active rule is matched against events and generates notifications. If the rule is already active, this command has no effect.

Access control

User Category	Permission

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- EVENT_RULE_NAME_DOES_NOT_EXIST
  The event rule name does not exist.

Creating event notification rules

Use the **rule_create** command to create an event notification rule.

```
rule_create rule=RuleName [ min_severity=<INFORMATIONAL|WARNING|MINOR|MAJOR|CRITICAL|NONE> ]
 [ codes=Codes | except_codes=EventCodes ] [ escalation_only=<yes|no> ]
 dests=dest1,dest2,... [ snooze_time=SnoozeTime ]
 [ escalation_time=EscalationTime escalation_rule=EscalationRule ] [ domain=DomainList ]
```

Parameters

Name	Type	Description	Mandatory	Default
rule	Object name	The name of the new rule.	Y	N/A
min_severity	Enumeration	Minimal event severity for rule filtering.	N	All severities.
codes	N/A	Filter only events with these codes.	N	All events.
except_codes	N/A	Filter only events with other codes.	N	All events.
escalation_only	Boolean	Specifies that this rule can only be used for escalation.	N	no
dests	Object name	Comma-separated list of destinations and destination groups for event notification.	Y	N/A
snooze_time	Integer	Snooze time in minutes.	N	No snoozing.
escalation_rule	Object name	Escalation rule.	N	N/A
escalation_time	Integer	Escalation time in minutes. Escalation time should not be smaller than snooze time. Refer to escalation_rule above for more information.	N	No escalation.
domain	N/A	The rule will be attached to the specified domains. To define more than one domain, separate them with a comma. To specify all existing domains, use "*".	N	none

This command defines a new event notification rule. An event notification rule determines which events should generate which notifications. When an event occurs, it is checked by all currently defined rules, based on which notifications are generated.

Each rule has a filtering and notification configuration.

The filtering configuration controls which events match this rule. The filtering can be based on the event's code, by specifying a minimum severity. When using this configuration, each event with a severity higher or equal to the rule's min_severity parameter matches this rule. Alternatively, the rule may match only a specific event code. Two filters can be combined for events whose severity depends on a run-time parameter.

The second part of a rule configuration is a list of destinations and destination groups that receive the notification when an event matches the filtering criteria. If a destination is included both in the rule and in one of the rule's destination groups, it still gets only one notification. The same applies if a destination is included in two destination groups, or if the event matches the filtering criteria of several rules, all using the same destination.

A rule can be defined as *alerting*, which means that notifications are sent repeatedly until the matching events are cleared using the Clearing alerting events command.

Clearing the event does not mean that the problem has been solved. It only means that it was noticed and there is no need to continue sending notifications.

The repeated sending of notifications can be defined by two ways:

- The snooze parameter causes the notifications to be sent again and again to the same destinations. The time in minutes between the repeated transmissions is determined by the **snooze** value.
- The escalation_time and escalation_rule parameters cause the notifications to be sent to the destination list of the escalation_rule if it is not cleared within **escalation_time** minutes.

Rules can escalate only to alerting rules (that is, to rules that have snooze or escalation definitions) in order to prevent a situation where notifications are stopped from being sent.

A rule cannot escalate to itself, nor can it be defined in a cyclic escalation of rules.

The escalation_only parameter defines a rule without filters, which can only be used as an escalation for other rules.

The snooze time cannot be greater than the escalation time.

It is not permitted to define new rules while there are uncleared alerting events.

The following example sends alerts upon critical events to John's cellular number and to the emails of all the IT staff. The alerts will be resent every 20 minutes until the events are cleared.

Example:

```
rule_create rule=critical_alerts min_severity=critical destinations=john-cell,itstaff snooze_time=20
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- EVENT_RULE_MAX_REACHED
  The maximum allowed number of event rules is already reached.
- EVENT_RULE_CANNOT_ESCALATE_TO_NON_ALERTING_RULES
  An event rule cannot be escalated to a non-alerting rule.

  Troubleshooting: An alerting rule can only be escalated to another escalation rule.
- DEST_APPEARS_TWICE
  The destination or destination group appears twice.
- EVENT_RULE_NAME_ALREADY_EXISTS
  The event rule name already exists.
- EVENT_RULE_NAME_DOES_NOT_EXIST
  The event rule name does not exist.
- NAME_IS_NEITHER_DEST_NOR_GROUP
  The specified ame is neither a destination group name nor a destination name.
- ESCALATION_TIME_MUST_BE_LARGER_THAN_SNOOZE_TIME
  Escalation time must be larger than snooze time.
- RULE_MAX_DESTS_REACHED
  The maximum allowed number of destinations and destination groups in a rule is already reached.
- EVENT_RULE_MUST_HAVE_FILTER
  An alerting event rule must have a filter represented by an event code or severity.
- EVENT_RULE_CANNOT_REFER_TO_INTERNAL_EVENT_CODES
  A user event rule cannot refer to internal event codes.
- ESCALATION_EVENT_RULE_CANNOT_HAVE_FILTER
  An escalation-only event rule cannot have code or min_severity specification.
- ESCALATION_EVENT_RULE_MUST_BE_ALERTING
  An escalation-only event rule must be an alerting rule.
- TOO_MANY_EVENT_CODES
  A maximum of *Maximum* return codes can be defined.
- EVENT_CODE_APPEARS_TWICE
  The return code '*Code*' appears twice in the list.

  Troubleshooting: Make sure that each return code appears in the list only once.
- UNRECOGNIZED_EVENT_CODE
  '*String*' is not a recognized return code.

  Troubleshooting: Consult the manual for the list of valid return codes.
- EVENT_RULE_CANNOT_HAVE_A_CATEGORY
  A user event rule cannot have a category definition.
- DOMAIN_DOESNT_EXIST
  The domain does not exist.

- DESTINATION_IS_NOT_IN_RULE_DOMAINS
  The destination must be included in rule domains.
- DESTGROUP_IS_NOT_IN_RULE_DOMAINS
  The destination groups must be included in rule domains.
- ESCALATION_RULE_NOT_IN_RULE_DOMAINS
  An escalation rule must belong to rule domains.
- EVENT_RULE_MUST_NOT_HAVE_SNMP_DEST
  According to the current system configuration state (snmp_type = NONE), an event rule must not have an SNMP destination.

Deactivating a rule

Use the **rule_deactivate** command to deactivate an event notification rule.

```
rule_deactivate rule=RuleName
```

Parameters

Name	Type	Description	Mandatory
rule	Object name	The name of the rule to be deactivated.	Y

A deactivated rule is not matched against events and does not generate notifications. If the rule is already inactive, then this command has no effect.

Inactive rules cannot be used as escalation rules.

The rules of type `escalation_only` cannot be deactivated.

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- EVENT_RULE_NAME_DOES_NOT_EXIST
  The event rule name does not exist.
- ESCALATION_ONLY_RULES_ARE_ALWAYS_ACTIVE
  An escalation-only event rule cannot be deactivated or activated.

Deleting event notification rules

Use the **rule_delete** command to delete an event notification rule.

```
rule_delete rule=RuleName
```

Parameters

Name	Type	Description	Mandatory
rule	Object name	The rule to be deleted.	Y

Rules that are defined as the escalation of other rules cannot be deleted.

It is not permitted to delete a rule while there are uncleared alerting events.

Example:

```
rule_delete rule=emergency_alerts
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_DELETE_RULE
  Are you sure you want to delete rule *Rule*?

Return codes

- CANNOT_CHANGE_EVENT_CONF_WITH_ALERTING_EVENTS
  Cannot change an event configuration while there are alerting events.

  Troubleshooting: Clear all alerting events before changing an event configuration.
- EVENT_RULE_NAME_DOES_NOT_EXIST
  The event rule name does not exist.
- EVENT_RULE_USED_FOR_ESCALATION_CAN_NOT_BE_DELETED
  The event rule is an escalation rule of another event rule. Therefore, it cannot be deleted.

  Troubleshooting: Delete all escalation rules that refer to this rule as their escalation rule.

Listing event notification rules

Use the **rule_list** command to list event notification rules.

```
rule_list [ rule=RuleName ] [ internal=<yes|no> ] [ domain=DomainName ]
```

Parameters

Name	Type	Description	Mandatory	Default
rule	Object name	The rule to be listed.	N	All rules.
internal	Enumeration	Filters XIV internal rules.	N	no
domain	Object name	The domain name.	N	All Domains

Example:

```
rule_list
```

Output:

```
Name             Minimum Severity  Event Code  Destinations
---------------  ----------------  ----------  ----------------
emergency_alerts  critical          all         john-cell,itstaff
```

Field ID	Field output	Default position
name	Name	1
min_severity	Minimum Severity	2
codes	Event Codes	3
except_codes	Except Codes	4
dests	Destinations	5
active	Active	6
escalation_time	Escalation Time	N/A
snooze_time	Snooze Time	N/A
escalation_rule	Escalation Rule	N/A
escalation_only	Escalation Only	7
category	Category	N/A
creator	Creator	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed

User Category	Permission
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Renaming event notification rules

Use the **rule_rename** command to rename an event notification rule.

```
rule_rename rule=RuleName new_name=Name
```

Parameters

Name	Type	Description	Mandatory
rule	Object name	The rule to be renamed.	Y
new_name	Object name	The new name of the rule.	Y

Example:

```
rule_rename rule=critical_alerts new_name=emergency_alerts
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- CANNOT_CHANGE_EVENT_CONF_WITH_ALERTING_EVENTS
  Cannot change an event configuration while there are alerting events.

  Troubleshooting: Clear all alerting events before changing an event configuration.
- EVENT_RULE_NAME_ALREADY_EXISTS
  The event rule name already exists.
- EVENT_RULE_NAME_DOES_NOT_EXIST
  The event rule name does not exist.

Updating an event notification rule

Use the **rule_update** command to update an event notification rule.

```
rule_update rule=RuleName [ min_severity=<INFORMATIONAL|WARNING|MINOR|MAJOR|CRITICAL|NONE> ] [ codes=Codes ]
 [ except_codes=EventCodes ] [ escalation_only=<yes|no> ] [ dests=dest1,dest2,... ]
 [ snooze_time=SnoozeTime ] [ escalation_time=EscalationTime ]
 [ escalation_rule=EscalationRule ] [ domain=DomainList ]
```

Parameters

Name	Type	Description	Mandatory	Default
rule	Object name	The name of the rule.	Y	N/A
min_severity	Enumeration	Minimum event severity for rule filtering.	N	Leave unchanged.
codes	N/A	Filter only events with this code.	N	Leave unchanged.
except_codes	N/A	Filter only events with other codes.	N	Leave unchanged.
escalation_only	Boolean	Specifies that this rule can only be used for escalation.	N	no

Name	Type	Description	Mandatory	Default
dests	Object name	Comma-separated list of destinations and destination groups for event notification.	N	Leave unchanged.
snooze_time	Integer	Snooze time in minutes.	N	Leave unchanged.
escalation_time	Integer	Escalation time in minutes.	N	Leave unchanged.
escalation_rule	Object name	Escalation rule.	N	Leave unchanged.
domain	N/A	The rule will be attached to the specified domains. To specify more than one domain, separate them with a comma. To specify all existing domains, use "*".	N	Leave unchanged.

This command updates the configuration of an event notification rule. All parameters and their descriptions are identical to the Creating event notification rules command.

Parameters which are not specified are not changed.

Example:

```
rule_update rule=critical_alerts min_severity=critical destinations=john-cell,itstaff snooze_time=30
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- EVENT_RULE_NAME_DOES_NOT_EXIST
  The event rule name does not exist.
- CANNOT_CHANGE_EVENT_CONF_WITH_ALERTING_EVENTS
  Cannot change an event configuration while there are alerting events.

  Troubleshooting: Clear all alerting events before changing an event configuration.
- EVENT_RULE_CANNOT_ESCALATE_TO_ITSELF
  An event rule cannot be its own escalation rule.
- EVENT_RULE_CANNOT_ESCALATE_TO_NON_ALERTING_RULES
  An event rule cannot be escalated to a non-alerting rule.

  Troubleshooting: An alerting rule can only be escalated to another escalation rule.
- DEST_APPEARS_TWICE
  The destination or destination group appears twice.
- EVENT_RULE_MISSING_ESCALATION_RULE
  An alerting event rule must have an escalation rule.

  Troubleshooting: If escalation time is specified, then an escalation rule must be specified as well.
- EVENT_RULE_MISSING_ESCALATION_TIME
  An alerting event rule must have escalation time.

  Troubleshooting: If an escalation rule is specified, then escalation time must be specified also.
- NAME_IS_NEITHER_DEST_NOR_GROUP
  The specified ame is neither a destination group name nor a destination name.
- ESCALATION_TIME_MUST_BE_LARGER_THAN_SNOOZE_TIME
  Escalation time must be larger than snooze time.
- RULE_MAX_DESTS_REACHED
  The maximum allowed number of destinations and destination groups in a rule is already reached.
- EVENT_RULE_MUST_HAVE_FILTER
  An alerting event rule must have a filter represented by an event code or severity.
- CYCLIC_ESCALATION_RULES_DEFINITION

Event rule escalation cannot be cyclic.

- EVENT_RULE_USED_FOR_ESCALATION_MUST_BE_ALERTING
  The event rule is an escalation rule of another event rule, and thus must be an alerting rule.
- EVENT_RULE_CANNOT_REFER_TO_INTERNAL_EVENT_CODES
  A user event rule cannot refer to internal event codes.
- ESCALATION_EVENT_RULE_CANNOT_HAVE_FILTER
  An escalation-only event rule cannot have code or min_severity specification.
- EVENT_RULE_CANNOT_HAVE_A_CATEGORY
  A user event rule cannot have a category definition.
- EVENT_RULE_CANNOT_HAVE_BOTH_CODES_AND_EXCEPTION_CODES
  An event rule cannot have both codes and exception codes.
- ESCALATION_EVENT_RULE_MUST_BE_ALERTING
  An escalation-only event rule must be an alerting rule.
- TOO_MANY_EVENT_CODES
  A maximum of *Maximum* return codes can be defined.
- EVENT_CODE_APPEARS_TWICE
  The return code '*Code*' appears twice in the list.

  Troubleshooting: Make sure that each return code appears in the list only once.
- UNRECOGNIZED_EVENT_CODE
  '*String*' is not a recognized return code.

  Troubleshooting: Consult the manual for the list of valid return codes.
- DOMAIN_DOESNT_EXIST
  The domain does not exist.
- DESTINATION_IS_NOT_IN_RULE_DOMAINS
  The destination must be included in rule domains.
- DESTGROUP_IS_NOT_IN_RULE_DOMAINS
  The destination groups must be included in rule domains.
- ESCALATION_RULE_NOT_IN_RULE_DOMAINS
  An escalation rule must belong to rule domains.
- EVENT_RULE_MUST_NOT_HAVE_SNMP_DEST
  According to the current system configuration state (snmp_type = NONE), an event rule must not have an SNMP destination.

Defining an SMS gateway

Use the **smsgw_define** command to define an SMS gateway.

```
smsgw_define smsgw=SMSGatewayName email_address=email subject_line=SubjectLineScheme email_body=EmailBodySchem
e [ smtpgw=<SMTPGW1[,SMTPGW2]...|ALL> ]
```

Parameters

Name	Type	Description	Mandatory	Default
smsgw	Object name	SMS gateway name.	Y	N/A
email_address	Token String	Format for the email address.	Y	N/A
subject_line	Token String	Format for the subject line.	Y	N/A
email_body	Token String	Format for the email body.	Y	N/A
smtpgw	Object name	List of SMTP gateways to be used.	N	The SMTP gateways defined in the smtpgw_prioritize command.

SMS gateways are used to send event notifications via SMS messages. SMS messages are sent via SMS-to-email servers. To definw a new SMS gateway, it is necessary to know how SMS messages are encapsulated in the email message.

When the system sends an SMS message, it uses the actual message text that describes the event and the destination number. The destination number is comprised from an area code and the local number. Both are specified when a destination is defined as described in the Defining a new event notification destination command.

The message text and the destination numbers can be embedded into the email message into various parts of the email message: destination address, subject line, or email body. This command defines how email messages are formatted, and how they contain the information of the specific SMS.

When defining an SMS gateway, three parameters must be specified in order to define the formatting:

- email_address: This is the email address used for sending the SMS via the email-to-SMS gateway.
- subject_line: This is the subject line of the outgoing email that will be converted to an SMS.
- email_body: This is the body of the outgoing email that will be converted to an SMS.

For each of these parameters, the value can be either fixed text, or an event text, or the destination phone number. This is done by embedding the following escape sequences into the text:

- {areacode}. This escape sequence is replaced by the destination's cellular number area code.
- {number}. This escape sequence is replaced by the destination's cellular local number.
- {message}. This escape sequence is replaced by the text to be shown to the user.
- \{, \}, \\. These are replaced by the {, } or \ respectively.

By default, the email to the email-to-SMS server is sent through the defined SMTP servers, prioritized by the Prioritizing SMTP gateways command. If needed, the user may define a specific SMTP gateway or gateways to be used for sending email to this email-to-SMS gateway.

Several SMS gateways can be defined. The system will try the SMS gateways, in the order specified in Prioritizing SMS gateways until it successfully connects to one of them. It is possible to define that specific SMS destinations will use specific SMS gateways (see Defining a new event notification destination).

Example:

```
smsgw_define smsgw=SMSGW1 email_address={areacode}{number}@sms2emailserver.example.com subject_line=SMS email_
body={message}
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- GATEWAY_MAX_REACHED
  The maximum allowed number of gateways is already reached.
- SMSGW_CANNOT_BE_DEFINED_WITHOUT_SMTPGW
  The SMS gateway cannot be defined if no SMTP gateway is defined.
- GATEWAY_NAME_DOES_NOT_EXIST
  The gateway name does not exist.
- GATEWAY_NAME_APPEARS_TWICE
  The gateway name appears twice on the list.
- GATEWAY_NAME_ALREADY_EXISTS
  The gateway name already exists.

Deleting an SMS gateway

Use the **smsgw_delete** command to delete an SMS gateway.

```
smsgw_delete smsgw=SMSGatewayName
```

Parameters

Name	Type	Description	Mandatory
smsgw	Object name	SMS gateway to be deleted.	Y

A gateway cannot be deleted if it is part of a notification rule or if it is being used by a destination.

Before deleting an SMS gateway, make sure that all alerting events are cleared.

Example:

```
smsgw_delete smsgw=external-SMSGW
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_DELETE_SMS_GATEWAY
  Are you sure you want to delete SMS gateway *Gateway*?

Return codes

- CANNOT_CHANGE_EVENT_CONF_WITH_ALERTING_EVENTS
  Cannot change an event configuration while there are alerting events.

  Troubleshooting: Clear all alerting events before changing an event configuration.
- GATEWAY_NAME_DOES_NOT_EXIST
  The gateway name does not exist.
- GATEWAY_USED_BY_DESTINATION
  The gateway is used by a destination.

Listing SMS gateways

Use the **smsgw_list** command to list SMS gateways.

```
smsgw_list [ smsgw=SMSGatewayName ]
```

Parameters

Name	Type	Description	Mandatory	Default
smsgw	Object name	Name of SMS gateway to list.	N	All gateways.

The command lists all SMS gateways, or a specific one. For each SMS gateway, all of its configuration information is listed.

Field ID	Field output	Default position
name	Name	1
email_address	Email Address	2
gateways	SMTP Gateways	3
subject_line	Subject Line	N/A
email_body	Email Body	N/A
priority	Priority	N/A

Example:

```
smsgw_list
```

Output:

```
Name    Email Address                                    SMTP Gateways
SMSGW1  {areacode}{number}@sms2emailserver.example.com   all
SMSGW2  {areacode}{number}@example.com                   all
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Prioritizing SMS gateways

Use the **smsgw_prioritize** command to set the priorities of the SMS gateways for sending SMS messages.

```
smsgw_prioritize order=<gw1[,gw2]...>
```

Parameters

Name	Type	Description	Mandatory
order	Object name	List of all SMS gateways ordered by priority.	Y

SMS messages can be sent to cell phones through one of the email-to-SMS gateways in this list. This command determines the order in which the storage system attempts to use these SMS gateways.

Only one gateway is used and subsequent gateways are only tried if the preceding ones in this priority list return an error.

Specific SMS destinations may define their own SMS gateways to be used when sending SMS to these destinations, regardless of this list.

Example:

```
smsgw_prioritize order=SMSGW1,SMSGW2
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- CANNOT_CHANGE_EVENT_CONF_WITH_ALERTING_EVENTS
  Cannot change an event configuration while there are alerting events.

  Troubleshooting: Clear all alerting events before changing an event configuration.
- GATEWAY_NAME_APPEARS_TWICE
  The gateway name appears twice on the list.
- GATEWAY_NAME_DOES_NOT_EXIST
  The gateway name does not exist.
- GATEWAY_NAME_MISSING_FROM_LIST
  The gateway name is missing from the list.

Renaming an SMS gateway

Use the **smsgw_rename** command to rename an SMS gateway.

```
smsgw_rename smsgw=SMSGatewayName new_name=Name
```

Parameters

Name	Type	Description	Mandatory
smsgw	Object name	SMS gateway to be renamed.	Y
new_name	Object name	New name for the SMS gateway.	Y

Before renaming an SMS gateway, make sure that all alerting events are cleared.

Example:

```
smsgw_rename smsgw=SMSGW2 new_name=external-SMSGW
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- CANNOT_CHANGE_EVENT_CONF_WITH_ALERTING_EVENTS
  Cannot change an event configuration while there are alerting events.

  Troubleshooting: Clear all alerting events before changing an event configuration.
- GATEWAY_NAME_ALREADY_EXISTS
  The gateway name already exists.
- GATEWAY_NAME_DOES_NOT_EXIST
  The gateway name does not exist.

Updating an SMS gateway

Use the **smsgw_update** command to update an SMS gateway.

```
smsgw_update smsgw=SMSGatewayName [ email_address=email ] [ subject_line=SubjectLineScheme ]
 [ email_body=EmailBodyScheme ] [ smtpgw=<SMTPGW1[,SMTPGW2]...|ALL> ]
```

Parameters

Name	Type	Description	Mandatory	Default
smsgw	Object name	SMS gateway name.	Y	N/A
email_address	Token String	Format for email address.	N	Leave unchanged.
subject_line	Token String	Format for subject line.	N	Leave unchanged.
email_body	Token String	Format for the email's body.	N	Leave unchanged.
smtpgw	Object name	List of SMTP gateways to be used.	N	The SMTP gateways defined in the smtpgw_prioritize command.

This command updates the configuration information of an existing SMS gateway. For the exact description and documentation of each parameter, see the documentation of Defining an SMS gateway.

This command cannot be executed while there are uncleared alerting events.

Parameters that are not specified will not be changed.

Example:

```
smsgw_update smsgw=SMSGW1 email_address={areacode}{number}@sms2emailserver.example.com subject_line=NextraSMS
email_body={message}
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- CANNOT_CHANGE_EVENT_CONF_WITH_ALERTING_EVENTS
  Cannot change an event configuration while there are alerting events.

  Troubleshooting: Clear all alerting events before changing an event configuration.
- GATEWAY_NAME_APPEARS_TWICE
  The gateway name appears twice on the list.

- GATEWAY_NAME_DOES_NOT_EXIST
  The gateway name does not exist.

Defining a new SMTP gateway

Use the **smtpgw_define** command to define an SMTP gateway.

```
smtpgw_define smtpgw=SMTPGatewayName address=Address [ from_address=<email|DEFAULT> ]
 [ reply_to_address=<email|DEFAULT> ]
```

Parameters

Name	Type	Description	Mandatory	Default
smtpgw	Object name	SMTP gateway name.	Y	N/A
address	N/A	SMTP gateway address (IP or DNS name).	Y	N/A
from_address	N/A	Sender's email address used for outgoing emails sent through this SMTP server.	N	DEFAULT (system-wide sender's address that applies to all servers).
reply_to_address	N/A	The reply to address used for outgoing emails sent through this SMTP server.	N	DEFAULT (system-wide reply-to address that applies to all servers).

Several email gateways can be defined to enable notification of events by email or sending SMS messages via email-to-SMS gateways. By default, the system attempts to send each email notification through the first gateway according to the order that you specify. Subsequent gateways are only tried if the first in line returns an error. A specific email destination, or a specific SMS gateway may be defined to use only specific SMTP gateways.

The SMTP protocol dictates that every email message must specify the email address of the sender. This sender address must be a valid address for two reasons:

- Many SMTP gateways require a valid sender address, otherwise they will not forward the email, as a security measure in order to prevent unauthorized usage of the SMTP server. Often this sender address must be limited to a specific domain.
- The sender's address is used as the destination for error messages generated by the SMTP gateways, such as: incorrect email address, full email mailbox and so on.

If the sender's address is not specified for a specific SMTP gateway, a global system-wide sender's address specified in Setting configuration parameters is used.

The user can also configure a reply-to address which is different from the sender's address, if it is required that the return emails be sent to another destination.

Example:

```
smtpgw_define smtpgw=mailserver1 address=smtp.example.com from_address=nextra@example.com reply_to_address=ne
xtraerrors@example.com
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- FROM_ADDRESS_NOT_DEFINED
  Neither the gateway's From Address nor the default From Address is defined.
- GATEWAY_MAX_REACHED
  The maximum allowed number of gateways is already reached.
- GATEWAY_NAME_ALREADY_EXISTS
  The gateway name already exists.

Deleting an SMTP gateway

Use the **smtpgw_delete** command to delete the specified SMTP gateway.

```
smtpgw_delete smtpgw=SMTPGatewayName
```

Parameters

Name	Type	Description	Mandatory
smtpgw	Object name	SMTP gateway to be deleted.	Y

A gateway cannot be deleted if it is part of a notification rule, is being used as an SMS gateway, or if it belongs to a destination.

Before deleting an SMTP gateway, make sure that all alerting events are cleared.

Example:

```
smtpgw_delete smtpgw=mailserverbackup
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_DELETE_SMTP_GATEWAY
  Are you sure you want to delete SMTP gateway *Gateway*?

Return codes

- CANNOT_CHANGE_EVENT_CONF_WITH_ALERTING_EVENTS
  Cannot change an event configuration while there are alerting events.

  Troubleshooting: Clear all alerting events before changing an event configuration.
- GATEWAY_NAME_DOES_NOT_EXIST
  The gateway name does not exist.
- GATEWAY_USED_BY_DESTINATION
  The gateway is used by a destination.
- GATEWAY_USED_BY_SMS_GATEWAY
  The gateway is used by an SMS Gateway.

Listing SMTP gateways

Use the **smtpgw_list** command to list SMTP gateways.

```
smtpgw_list [ smtpgw=SMTPGatewayName ]  [ internal=<yes|no> ]
```

Parameters

Name	Type	Description	Mandatory	Default
smtpgw	Object name	Name of SMTP gateway to list.	N	no.
internal	Enumeration	Filters gateways by their XIV-internal attribute.	N	no

This command lists defined SMTP gateways and their configuration information.

Field ID	Field output	Default position
name	Name	1
address	Address	2
priority	Priority	3
from_address	From Address	N/A
reply_to_address	Reply-to Address	N/A
failed	Failed	N/A
port	Port	N/A
creator	Creator	N/A

Example:

```
smtpgw_list
```

Output:

```
Name          Email Address          Port  Priority
-----------   --------------------   ----  --------
mailserver1   smtp.example.com       25    1
mailserver2   smtp.example.com       25    2
```

Access control

User Category	Permission	Condition
Storage administrator	Conditionally Allowed	Allowed, unless the internal parameter is specified.
Storage integration administrator	Disallowed	N/A
Application administrator	Conditionally Allowed	Allowed, unless the internal parameter is specified.
Security administrator	Disallowed	N/A
Read-only users	Conditionally Allowed	Allowed, unless the internal parameter is specified.
Technicians	Allowed	N/A

Prioritizing SMTP gateways

Use the **smtpgw_prioritize** command to prirotize SMTP gateways.

```
smtpgw_prioritize order=<gw1[,gw2]...>
```

Parameters

Name	Type	Description	Mandatory
order	Object name	List of all the SMTP gateways in the order of their priority.	Y

Several email gateways can be defined to enable notification of events or the sending of SMS by email. By default, the storage system attempts to send each email through the first gateway according to the order that is specified in this command. Only one gateway is used and subsequent gateways are only tried if the preceding ones in this priority list return an error.

These priorities are used only for email destinations and SMS gateways that did not specify their own SMTP gateways.

Example:

```
smtpgw_prioritize order=mailserver2,mailserver1
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- CANNOT_CHANGE_EVENT_CONF_WITH_ALERTING_EVENTS
  Cannot change an event configuration while there are alerting events.

  Troubleshooting: Clear all alerting events before changing an event configuration.
- GATEWAY_NAME_APPEARS_TWICE
  The gateway name appears twice on the list.
- GATEWAY_NAME_DOES_NOT_EXIST
  The gateway name does not exist.
- GATEWAY_NAME_MISSING_FROM_LIST
  The gateway name is missing from the list.

Renaming an SMTP gateway

Use the **smtpgw_rename** command to rename an SMTP gateway.

```
smtpgw_rename smtpgw=SMTPGatewayName new_name=Name
```

Parameters

Name	Type	Description	Mandatory
smtpgw	Object name	SMTP gateway to be renamed.	Y
new_name	Object name	New name for the SMTP gateway.	Y

Example:

```
smtpgw_rename smtpgw=mailserver2 new_name=mailserverbackup
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- CANNOT_CHANGE_EVENT_CONF_WITH_ALERTING_EVENTS
  Cannot change an event configuration while there are alerting events.

  Troubleshooting: Clear all alerting events before changing an event configuration.
- GATEWAY_NAME_ALREADY_EXISTS
  The gateway name already exists.
- GATEWAY_NAME_DOES_NOT_EXIST
  The gateway name does not exist.

Updating an SMTP gateway

Use the **smtpgw_update** command to update the configuration of an SMTP gateway.

```
smtpgw_update smtpgw=SMTPGatewayName [ address=Address ] [ from_address=<email|DEFAULT> ]
 [ reply_to_address=<email|DEFAULT> ] [ internal=<yes|no> ] [ port=PortNumber ]
```

Parameters

Name	Type	Description	Mandatory	Default
smtpgw	Object name	SMTP gateway name.	Y	N/A
address	N/A	SMTP gateway address (IP or DNS name).	N	Leave unchanged.
internal	Boolean	For an XIV internal gateway, set to Yes.	N	NO
from_address	N/A	Sender's email address used for out-going emails sent through this SMTP server, or DEFAULT for the system-wide default.	N	Leave unchanged.
reply_to_address	N/A	The reply-to address used for outgoing emails sent through this SMTP server, or DEFAULT for the system-wide default.	N	Leave unchanged.
port	Integer	TCP port used in the gateway instead of the default port 25.	N	Leave unchanged.

This command updates the configuration of an existing SMTP gateway. Fields which are not specified are not changed.

Example:

```
smtpgw_update smtpgw=mailserver1 address=smtp2.example.com from_address=nextra@example.com reply_to_address=ne
xtraerrors@example.com
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission	Condition
Storage administrator	Conditionally Allowed	Allowed, unless the internal parameter is specified.
Storage integration administrator	Disallowed	N/A
Application administrator	Conditionally Allowed	Allowed, unless the internal parameter is specified.
Security administrator	Disallowed	N/A
Read-only users	Disallowed	N/A
Technicians	Allowed	N/A

Return codes

- GATEWAY_NAME_DOES_NOT_EXIST
  The gateway name does not exist.
- CANNOT_CHANGE_EVENT_CONF_WITH_ALERTING_EVENTS
  Cannot change an event configuration while there are alerting events.

  Troubleshooting: Clear all alerting events before changing an event configuration.
- FROM_ADDRESS_NOT_DEFINED
  Neither the gateway's From Address nor the default From Address is defined.

Generating an XMPNS admin control event

Use the **xmpns_admin_config_set** command to generate an XMPNS_ADMIN_CONTROL event.

```
xmpns_admin_config_set action=Action user=User
```

Parameters

Name	Type	Description	Mandatory
action	String	Action code text.	Y
user	String	User name.	Y

This command generates an XMPNS_ADMIN_CONTROL event which includes the **action_codec** text in the event's description field. The username is also added to the action string sent in the description field.

Example:

```
xmpns_admin_config_set action user
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

IP configuration commands

This section describes the command-line interface (CLI) for IP configuration.

- **Updating the configuration of the Ethernet port**
  Use the **ethernet_port_update** command to update the configuration of the Ethernet port.
- **Showing the status and configuration of Ethernet ports (deprecated)**
  Use the **ipinterface_list_ports** command to list all Ethernet ports together with their configurations and statuses.
- **Showing the status and configuration of Ethernet ports**
  Use the **ethernet_port_list** command to list all Ethernet ports together with their configurations and statuses.
- **Displaying Ethernet (physical) port statistics in the system**
  Use the **ethernet_port_counter_list** command to display Ethernet (physical) port statistics in the system.
- **Creating a new IP interface**
  Use the **ipinterface_create** command to create a new IP interface for iSCSI ports.
- **Deleting IP interfaces**
  Use the **ipinterface_delete** command to delete an IP interface.

- **Listing IP interface configuration**
  Use the **ipinterface_list** command to list the configuration of a specific IP interface or all IP interfaces.
- **Listing IP interface addresses**
  Use the **ipinterface_list_ips** command to list the IP addresses configured on a specific IP interface or all IP interfaces.
- **Displaying IP interface statistics in the system**
  Use the **ipinterface_counter_list** command to display the IP interfaces and their statistics in the system.
- **Renaming an IP interface**
  Use the **ipinterface_rename** command to rename an IP interface.
- **Printing the ARP database of an IP interface**
  Use the **ipinterface_run_arp** command to print the ARP database of the specified IP interface.
- **Testing the traceroute to a remote IP**
  Use the **ipinterface_run_traceroute** to test connectivity to a remote IP node using the ICMP trace-route mechanism.
- **Testing the traceroute to a remote IPv6 address**
  Use the **ipinterface_run_traceroute6** command to test connectivity to a remote IPv6 address using the ICMP trace-route mechanism.
- **Updating an IP interface**
  Use the **ipinterface_update** command to update the configuration of an IP interface.
- **Updating VLAN priority code points**
  Use the **ipinterface_vlan_update** command to update the priority code points of all IP interfaces which have a VLAN ID defined.
- **Defining a new IPSec connection**
  Use the **ipsec_connection_add** command to add a new IPSec connection.
- **Updating an existing IPSec connection**
  Use the **ipsec_connection_update** command to update an existing IPSec connection.
- **Removing an existing IPSec connection**
  Use the **ipsec_connection_remove** command to remove an existing IPSec connection.
- **Listing IPSec connections**
  Use the **ipsec_connection_list** command to list all or specific IPSec connections.
- **Listing IPSec tunnels**
  Use the **ipsec_list_tunnels** command to list IPSec tunnels.
- **Connecting to a support center**
  Use the **support_center_connect** command to connect to a support center.
- **Defining a support center**
  Use the **support_center_define** command to define a support center.
- **Deleting a support center**
  Use the **support_center_delete** command to delete a support center.
- **Disconnecting from a support center**
  Use the **support_center_disconnect** command to disconnect the storage system from a support center.
- **Listing support centers**
  Use the **support_center_list** command to list support centers.
- **Listing statuses of all support centers**
  Use the **support_center_status** command to list information about the status of the defined support centers.
- **Configuring the support center connection to enable automatic connect on restart**
  Use the **support_center_config** command to configure the connection to a support center automatically.
- **Listing the configuration of support center automatically connect**
  Use the **support_center_config_list** command to display the configuration of the automatic connection to a support center.
- **Creating a new IP access group**
  Use the **ip_access_group_create** command to create a new IP access group.
- **Removing an address from an IP access group**
  Use the **ip_access_group_remove_address** command to delete the IP address of an access group.
- **Adding a new address to an IP access group**
  Use the **ip_access_group_add_address** command to add a new IP to an access group.
- **Deleting an existing IP access group**
  Use the **ip_access_group_delete** command to delete an IP access group.
- **Renaming an existing IP access group**
  Use the **ip_access_group_rename** command to rename an existing IP access group.
- **Listing IP access groups**
  Use the **ip_access_group_list** command to list IP access groups.
- **Listing IP access group addresses**
  Use the **ip_access_group_address_list** command to list IP access group addresses.

Updating the configuration of the Ethernet port

Use the **ethernet_port_update** command to update the configuration of the Ethernet port.

```
ethernet_port_update ethernet_port_name=EthernetPortName mtu=MTU [ force_mtu_change=<yes|no> ]
```

Parameters

Name	Type	Description	Mandatory	Default
ethernet_port_name	Object name	Name of the Ethernet port. The name format is ethernet_m_p. 'm' is the module_id and 'p' is the port.	Y	N/A
mtu	Integer	Maximum Transmission Unit: The packet size that is supported by the connecting Ethernet switch. Valid values 1500 - 9000.	Y	N/A

Name	Type	Description	Mandatory	Default
force_mtu_change	Boolean	If yes, forces the MTU changes for management or VPN Ethernet port.	N	no

This command updates the configuration of the Ethernet port.

Example:

```
ethernet_port_update ethernet_port_name=ethernet_port_4_5 mtu=2000
```

Output:

```
Command completed successfully
```

Attention:
Be careful when changing Ethernet port MTU. If management connectivity is on, or a VPN is being used, this might break existing application connectivity, unless the adjacent switch is properly configured. This would be the result of the switch using a fixed MTU size lower than the new size being defined on the system.

For example, if the system sends 9000 byte packets while the switch can only receive packets up to 1500 bytes, the switch drops the extra packets.

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Allowed

Return codes

- ETHERNET_PORT_DOES_NOT_EXIST
  The Ethernet port does not exist.
- ETHERNET_PORT_MTU_UPDATE_NOT_SUPPORTED_FOR_MANAGEMENT_OR_VPN
  Ethernet port MTU update not supported for management or VPN ports.

  Troubleshooting: Changing Ethernet port MTU with management or VPN on, might break existing applications connectivity. Verify that the adjacent switch supports the newly configured MTU size.

Showing the status and configuration of Ethernet ports (deprecated)

Use the **ipinterface_list_ports** command to list all Ethernet ports together with their configurations and statuses.

```
ipinterface_list_ports
```

Note:
As of version 12.3.2, this is a legacy command. Instead, use the **ethernet_port_list** command.

All physical Ethernet ports used to connect to the user's network are listed. The list includes the following information:

- Component ID (Module number for iSCSI or switch number for management/field technician port)
- Port number on module/switch
- For management/VPN/field technician: "management"/"VPN"/"field technician"
- IP interface containing the ports (or none, if port is not configured as part of IP interface)
- Status up/down
- Auto-negotiation: Half-full duplex, 1000/100/10

Example:

```
ipinterface_list_ports
```

Output:

```
Index   Role          IP Interface   Connected Component    Link Up?
------- ------------ -------------- ---------------------- ----------
1       Component                    1:Flash_Canister:4:1   yes
1       Component                    1:Flash_Canister:4:2   yes
1       IPMI                         1:Module:13            yes
1       IPMI                         1:Module:14            yes
1       IPMI                         1:Module:9             yes
1       Internal                     1:IB_Switch:1:12       yes
1       Internal                     1:IB_Switch:1:13       yes
1       Internal                     1:IB_Switch:1:8        yes
1       Management                                          yes
1       iSCSI                                               unknown
1       iSCSI                                               unknown
```

```
1       iSCSI                                           unknown
2       IPMI                        1:Module:11         yes
2       IPMI                        1:Module:12         yes
2       IPMI                        1:Module:7          yes
2       iSCSI                                           unknown
2       iSCSI                                           unknown
2       iSCSI                                           unknown
```

```
Cont.:
```

```
Negotiated Speed (Mb/s)   Full Duplex?   Module         RX Flow Control?   TX Flow Control?
------------------------- -------------- -------------- ------------------ ------------------
1000                      yes            1:Module:12    yes                yes
1000                      yes            1:Module:13    yes                yes
1000                      yes            1:Module:12    yes                yes
1000                      yes            1:Module:13    yes                yes
1000                      yes            1:Module:8     yes                yes
10000                     yes            1:Module:12    yes                yes
10000                     yes            1:Module:13    yes                yes
10000                     yes            1:Module:8     yes                yes
1000                      yes            1:Module:12    yes                yes
N/A                       unknown        1:Module:12    yes                yes
N/A                       unknown        1:Module:13    yes                yes
N/A                       unknown        1:Module:8     yes                yes
1000                      yes            1:Module:12    yes                yes
1000                      yes            1:Module:13    yes                yes
1000                      yes            1:Module:8     yes                yes
N/A                       unknown        1:Module:12    yes                yes
N/A                       unknown        1:Module:13    yes                yes
N/A                       unknown        1:Module:8     yes                yes
```

Field ID	Field output	Default position
index	Index	1
role	Role	2
ip_interface_name	IP Interface	3
connected_component	Connected Component	4
is_link_up	Link Up?	5
negotiated_speed_Mbs	Negotiated Speed (Mb/s)	6
is_full_duplex	Full Duplex?	7
module_id	Module	8
requires_service	Requires Service	N/A
service_reason	Service Reason	N/A
pause_autonegotiate	Flow control auto-negotiate?	N/A
pause_rx	RX Flow Control?	9
pause_tx	TX Flow Control?	10

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Showing the status and configuration of Ethernet ports

Use the **ethernet_port_list** command to list all Ethernet ports together with their configurations and statuses.

```
ethernet_port_list
```

All physical Ethernet ports used to connect to the user's network are listed. The list includes the following information:

- Component ID (Module number for iSCSI or switch number for management/field technician port)
- Port number on module/switch
- For management/VPN/field technician: "management"/"VPN"/"field technician"
- IP interface containing the ports (or none, if port is not configured as part of IP interface)
- Status up/down
- Auto-negotiation: Half-full duplex, 1000/100/10

Example:

```
ethernet_port_list
```

Output:

```
Port Name            Role                  Connected Component    Link Up?
-------------------- --------------------- ---------------------- ----------
ethernet_port_1_1    IPMI_1                1:Module:2             yes
ethernet_port_1_2    IPMI_2                1:Module:3             yes
ethernet_port_1_3    Management                                   yes
ethernet_port_1_4    canister_eth_port_2   1:Flash_Canister:1:2   yes
ethernet_port_1_5    Data                                         yes
ethernet_port_1_6    Data                                         no
ethernet_port_2_1    IPMI_1                1:Module:3             yes
ethernet_port_2_2    IPMI_2                1:Module:1             yes
ethernet_port_2_3    Management                                   yes
ethernet_port_2_4    canister_eth_port_1   1:Flash_Canister:1:1   yes
ethernet_port_2_5    Data                                         yes
ethernet_port_2_6    Data                                         no
ethernet_port_3_1    IPMI_1                1:Module:1             yes
ethernet_port_3_2    IPMI_2                1:Module:2             yes
ethernet_port_3_3    Management                                   yes
ethernet_port_3_5    Data                                         yes
ethernet_port_3_6    Data                                         no

cont:

Negotiated Speed (Mb/s)     MTU    Full Duplex?   Module
--------------------------  ------ -------------- ------------
1000                        1500   yes            1:Module:1
1000                        1500   yes            1:Module:1
1000                        1500   yes            1:Module:1
1000                        1500   yes            1:Module:1
10000                       1500   yes            1:Module:1
0                           1500   no             1:Module:1
1000                        1500   yes            1:Module:2
1000                        1500   yes            1:Module:2
1000                        1500   yes            1:Module:2
1000                        1500   yes            1:Module:2
10000                       1500   yes            1:Module:2
0                           1500   no             1:Module:2
1000                        1500   yes            1:Module:3
1000                        1500   yes            1:Module:3
1000                        1500   yes            1:Module:3
10000                       1500   yes            1:Module:3
0                           1500   no             1:Module:3

cont:
Flow control auto-negotiate?   RX Flow Control?   TX Flow Control?
------------------------------ ------------------ ------------------
yes                            yes                yes
yes                            yes                yes
yes                            yes                yes
yes                            yes                yes
no                             yes                yes
no                             yes                yes
yes                            yes                yes
yes                            yes                yes
yes                            yes                yes
yes                            yes                yes
no                             yes                yes
no                             yes                yes
yes                            yes                yes
yes                            yes                yes
yes                            yes                yes
no                             yes                yes
no                             yes                yes
```

Field ID	Field output	Default position
ethernet_port_name	Port Name	1
role	Role	2
connected_component	Connected Component	3
is_link_up	Link Up?	4
negotiated_speed_Mbs	Negotiated Speed (Mb/s)	5
mtu	MTU	6
is_full_duplex	Full Duplex?	7
module_id	Module	8
requires_service	Requires Service	N/A
service_reason	Service Reason	N/A
pause_autonegotiate	Flow control auto-negotiate?	9
pause_rx	RX Flow Control?	10
pause_tx	TX Flow Control?	11
nic_component_id	NIC	N/A
cna_component_id	CNA	N/A

Access control

User Category	Permission

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Displaying Ethernet (physical) port statistics in the system

Use the **ethernet_port_counter_list** command to display Ethernet (physical) port statistics in the system.

```
ethernet_port_counter_list
```

Example:

```
ethernet_port_counter_list
```

Output:

```
Port Name            Role          Module       Connected Component
-------------------- ------------ ------------ ----------------------
ethernet_port_1_1    IPMI          1:Module:1   1:Module:2
ethernet_port_1_2    IPMI          1:Module:1   1:Module:3
ethernet_port_1_3    Management    1:Module:1
ethernet_port_1_4    Component     1:Module:1   1:Flash_Canister:1:2
ethernet_port_1_5    iSCSI         1:Module:1
ethernet_port_1_6    iSCSI         1:Module:1
ethernet_port_2_1    IPMI          1:Module:2   1:Module:3
ethernet_port_2_2    IPMI          1:Module:2   1:Module:1
ethernet_port_2_3    Management    1:Module:2
ethernet_port_2_4    Component     1:Module:2   1:Flash_Canister:1:1
ethernet_port_2_5    iSCSI         1:Module:2
ethernet_port_3_1    IPMI          1:Module:3   1:Module:1
ethernet_port_3_2    IPMI          1:Module:3   1:Module:2
ethernet_port_3_5    iSCSI         1:Module:3
ibtun0               Internal      1:Module:1   1:Module:2
ibtun0               Internal      1:Module:2   1:Module:3
ibtun0               Internal      1:Module:3   1:Module:1
```

```
cont:
RX bytes     RX packets   TX bytes     TX packets   RX errors   RX drops
------------ ------------ ------------ ------------ ----------- ----------
814582       8131         814582       8131         0           0
815590       8145         815860       8148         0           0
21626657     271999       2613884      5554         0           0
10044406     18337        3439149      23587        0           0
0            0            0            0            0           0
0            0            0            0            0           0
811622       8083         812154       8089         0           0
814916       8134         815092       8136         0           0
20182785     266284       29959        312          0           0
142025271    106676       7628675      92402        0           0
0            0            0            0            0           0
816288       8152         816370       8153         0           0
812154       8089         811892       8086         0           0
0            0            0            0            0           0
210368770    1587299      537221876    1287096      0           0
246042184    1898524      594705380    1026429      0           0
1099820788   2144701      260338771    1159386      0           0
```

Field ID	Field output	Default position
ethernet_port_name	Port Name	1
role	Role	2
connected_component	Connected Component	3
rx_bytes	RX bytes	4
rx_pkts	RX packets	5
tx_bytes	TX bytes	6
tx_pkts	TX packets	7
rx_errors	RX errors	8
rx_drop	RX drops	9
missing	missing	N/A

Access control

User Category	Permission
Storage administrator	Allowed

User Category	Permission
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Creating a new IP interface

Use the **ipinterface_create** command to create a new IP interface for iSCSI ports.

```
ipinterface_create ipinterface=IPInterfaceName address=Address netmask=NetworkMask [ gateway=DefaultGateway ]
 < < ethernet_port_name=EthernetPortName [ vlan_id=VlanID ]
 > | < module=ModuleNumber port=PortNumber [ mtu=MTU ]
 [ speed=<auto|10mb|100mb|1000mb|1gb|2500mb|2.5gb|10000mb|10gb> ] > >
```

Parameters

Name	Type	Description	Mandatory	Default
ipinterface	Object name	The name of the IP interface to be created. Do not use the names `Management` or `VPN`.	Y	N/A
address	N/A	IP address of the interface.	Y	N/A
netmask	N/A	Network mask of the interface.	Y	N/A
gateway	N/A	IP address of the default gateway for this interface.	N	None
ethernet_port_name	Object name	Name of the Ethernet port. The name format is ethernet_m_p. 'm' is the module_id and 'p' is the port.	N	N/A
vlan_id	Integer	VLAN id. Valid values 1-4094.	N	0
module	N/A	Component identifier (rack and module) of the module containing Ethernet ports.	N	N/A
port	Integer	Port Number.	N	N/A
mtu	Integer	Maximum Transmission Unit: The supported packet size by the connecting Ethernet switch. This is optional when the default equals 1536. MTU of up to 4500 is supported.	N	4500 for iSCSI and 1536 for Management and VPN.
speed	Enumeration	Interface's speed. An explicit speed turns off auto-negotiation. Valid values: 'auto' or 'explicit'.	N	auto

This command defines a new IP interface for iSCSI traffic. Gateway, MTU, network mask and IP are the standard IP definitions.

Each iSCSI Ethernet port can be defined as an IP interface.

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Allowed

Return codes

- IPINTERFACE_EXISTS
  This IP interface name is already in use.
- ILLEGAL_PORT_NUMBER
  The port number is out of range.
- ETHERNET_PORT_DOES_NOT_EXIST
  The Ethernet port does not exist.
- PORT_IS_USED_IN_ANOTHER_IP_INTERFACE
  One of the physical ports specified is already assigned to an IP Interface.
- IP_ADDRESS_ALREADY_USED_IN_ANOTHER_INTERFACE
  The IP address is already assigned to another interface.
- IPADDRESS_AND_GATEWAY_ARE_NOT_ON_SAME_SUBNET
  The IP address specified for the default gateway is not in the subnet of the IP interface.

- MTU_TOO_LARGE
  The specified MTU value is too large.
- ILLEGAL_COMPONENT_ID
  This component ID is illegal.
- ILLEGAL_IPADDRESS
  An illegal IP address was entered.
- DUPLICATE_IPADDRESSES
  Duplicate IP addresses were specified.
- ILLEGAL_GATEWAY_IPADDRESS
  An illegal IP address was specified for the default gateway.
- TOO_MANY_VLANS
  The number of VLAN IDs is larger than that allowed by the system.
- TOO_MANY_IPINTERFACES
  The number of IP interfaces is too large.
- TOO_MANY_IPINTERFACES_PER_ETHERNET_PORT
  The number of IP interfaces per Ethernet port is larger than that allowed by the system.
- TOO_MANY_IPINTERFACES_PER_VLAN
  The number of IP interfaces per VLAN is larger than that allowed by the system.
- IP_ADDRESS_ALREADY_EXISTS
  The system already has the same IP address on the same VLAN.
- ILLEGAL_ETHERNET_PORT_WITH_VLAN
  The Ethernet port role does not support VLANs. Only management, VPN, or iSCSI roles are allowed.
- TWO_IPINTERFACES_ON_SAME_ETHERNET_PORT_AND_VLAN_NOT_ALLOWED
  Two IP interfaces on the same Ethernet port and having the same VLAN ID are not allowed.
- MAX_IP_INTERFACES_FOR_TARGET_CONNECTIVITY_PER_ETHERNET_PORT_REACHED
  Maximal number of IP Interfaces for target connectivity using same port reached.
- PARAMETER_CANNOT_BE_UPDATED_ON_IPINTERFACE_WITH_VLAN
  The MTU, speed, module or port parameter can be used only if VLAN is not configured and this is the only IP interface on the physical port.

Deleting IP interfaces

Use the **ipinterface_delete** command to delete an IP interface.

```
ipinterface_delete ipinterface=IPInterfaceName
```

Parameters

Name	Type	Description	Mandatory
ipinterface	Object name	The IP interface to be deleted.	Y

Only the interfaces defined for iSCSI traffic can be deleted. Management and VPN interfaces cannot be deleted.

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Allowed

Return codes

- IPINTERFACE_DOES_NOT_EXIST
  This IP interface name does not exist.
- COMMAND_NOT_ALLOWED_ON_MANAGEMENT_OR_VPN_INTERFACE
  The operation is not allowed on the management or VPN IP Interface.
- IPINTERFACE_HAS_CONNECTIVITY
  The IP interface has connectivity defined to another machine.

Listing IP interface configuration

Use the **ipinterface_list** command to list the configuration of a specific IP interface or all IP interfaces.

```
ipinterface_list [ ipinterface=IPInterfaceName | address=Address | address6=IPv6address ]
```

Parameters

Name	Type	Description	Mandatory	Default
ipinterface	Object name	The IP interface to be listed.	N	All interfaces
address	N/A	IP address of the interface to be listed.	N	All interfaces
address6	N/A	IPv6 address of the interface to be listed.	N	All interfaces

This command lists configuration information for the specified IP interface, or for all IP interfaces (including management). The management or VPN name can only be used to view the configuration of the management of VPN interfaces.

The following information is listed:

- Name
- Type (iSCSI/management)
- IP address (or comma separated addresses for management and VPN)
- Network mask
- Default gateway
- CIDR address (or comma separated addresses for management and VPN)
- Default IPv6 gateway
- MTU
- Module (for iSCSI only)
- Comma separated list of ports (for iSCSI only)
- Interface desired speed information

Example:

```
ipinterface_list
```

Output:

```
Name          Type          IP Address      Network Mask    Default Gateway   IPv6 Address
------------ ------------ --------------- --------------- ----------------- --------------
management   Management   192.0.2.1       255.255.255.0   192.0.2.254
```

```
Cont.:

IPv6 Gateway   MTU    Module         Port   IP access group name
-------------- ------ -------------- ------ ----------------------
               1500   1:Module:12
```

Field ID	Field output	Default position
name	Name	1
ethernet_port_name	Physical Port Name	2
type	Type	3
vlan_id	VLAN ID	4
vlan_pcp	VLAN PCP	5
address	IP Address	6
netmask	Network Mask	7
gateway	Default Gateway	8
address6	IPv6 Address	9
gateway6	IPv6 Gateway	10
mtu	MTU	11
module	Module	12
port	Port	13
speed	Speed	N/A
access_group	IP access group name	14

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing IP interface addresses

Use the **ipinterface_list_ips** command to list the IP addresses configured on a specific IP interface or all IP interfaces.

```
ipinterface_list_ips [ ipinterface=IPInterfaceName | address=Address | address6=IPv6address | module=ModuleNum
ber ]
```

Parameters

Name	Type	Description	Mandatory	Default
ipinterface	Object name	The IP interface to be listed.	N	All interfaces
address	N/A	IP address of the interface to be listed.	N	All addresses
address6	N/A	IPv6 address of the interface to be listed.	N	All addresses
module	N/A	Limits the listing to a specific module.	N	All modules

This command lists IP addresses for the specified interface, or for the specified module, or for both (including Management). The Management or VPN name can only be used to view IP addresses configured for the management of VPN interfaces.

The following information is listed:

- IP Interface
- Interface Type (iSCSI/Management/VPN)
- Address (in CIDR format)
- Address type (Static IPv4/Static IPv6/Link Local IPv6/Site Local IPv6/Global IPv6)
- Module

Example:

```
ipinterface_list_ips
```

Output:

```
IP Interface   Interface Type   Address
-------------- ---------------- ------------------------------------------
management     Management       2001:ODB8::0001/32
management     Management       192.0.2.1/24
management     Management       2001:ODB8:42f2:e9ff:feaf:ccb2/32

Cont.:

Address Type      Module         IP access group name
----------------- -------------- ----------------------
Global IPv6       1:Module:12
Static IPv4       1:Module:12
Link Local IPv6   1:Module:12
```

Field ID	Field output	Default position
ipinterface	IP Interface	1
ipinterface_type	Interface Type	2
address	Address	3
address_type	Address Type	4
module	Module	5
access_group	IP access group name	6
vlan_id	VLAN ID	7

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Displaying IP interface statistics in the system

Use the **ipinterface_counter_list** command to display the IP interfaces and their statistics in the system.

```
ipinterface_counter_list
```

Example:

```
ipinterface_counter_list
```

Output:

```
Role          Name                   RX packets   TX packets   RX errors
------------ ---------------------- ------------ ------------ -----------
Management   ethernet_port_1_3      244553       5580         0
Management   ethernet_port_2_3      239749       324          0
Management   ethernet_port_3_3      239768       328          0
iSCSI        test_iscsi_m2_p2       0            0            0
iSCSI        test_old_iscsi_m2_p2   0            0            0
```

Field ID	Field output	Default position
role	Role	1
name	Name	2
rx_pkts	RX packets	3
tx_pkts	TX packets	4
rx_errors	RX errors	5

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Renaming an IP interface

Use the **ipinterface_rename** command to rename an IP interface.

```
ipinterface_rename ipinterface=IPInterfaceName new_name=Name
```

Parameters

Name	Type	Description	Mandatory
ipinterface	Object name	Original name of the IP interface.	Y
new_name	Object name	The new name of the IP interface.	Y

This command renames an IP interface. The IP interface must be unique in the system. This command cannot be applied to Management or VPN interfaces.

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Allowed

Return codes

- IPINTERFACE_DOES_NOT_EXIST
  This IP interface name does not exist.
- IPINTERFACE_EXISTS
  This IP interface name is already in use.
- COMMAND_NOT_ALLOWED_ON_MANAGEMENT_OR_VPN_INTERFACE
  The operation is not allowed on the management or VPN IP Interface.

Printing the ARP database of an IP interface

Use the **ipinterface_run_arp** command to print the ARP database of the specified IP interface.

```
ipinterface_run_arp < localipaddress=IPaddress [ vlan_id=VlanID ] > | localipaddress6=IPv6address
```

Parameters

Name	Type	Description	Mandatory	Default
localipaddress	N/A	IP address of the IP interface for the ARP database to be printed.	N	N/A
vlan_id	Integer	The virtual LAN ID. Valid values 1-4094.	N	0
localipaddress6	N/A	IPv6 address of the IP interface for the ARP database to be printed.	N	N/A

This command prints a list of the ARP database of an IP interface with its IP addresses and their associated Ethernet MAC addresses. The IP address must be one of the IP addresses defined for iSCSI IP interfaces, or the management or VPN name.

Field ID	Field output	Default position
arp_output	arp Output	1

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Allowed

Return codes

- NO_IP_INTERFACE_MATCHES_CRITERIA
  No IP Interface or VLAN ID matches the defined criteria. Check if the IP Address and/or VLAN ID are correct.
- MORE_THAN_ONE_IP_INTERFACE_MATCHES
  More than one IP Interface matches the defined criteria.

Testing the traceroute to a remote IP

Use the **ipinterface_run_traceroute** to test connectivity to a remote IP node using the ICMP trace-route mechanism.

```
ipinterface_run_traceroute localipaddress=IPaddress remote=remoteHost [ vlan_id=VlanID ]
```

Parameters

Name	Type	Description	Mandatory	Default
localipaddress	N/A	IP address of the IP interface for which the traceroute command is run.	Y	N/A
remote	N/A	IP address or DNS for the traceroute test.	Y	N/A
vlan_id	Integer	The virtual LAN ID. Valid values 1-4094.	N	0

This command runs a route trace to the specified remote host through the specified IP interface. The IP address must be one of the IP addresses defined for iSCSI IP interfaces or the Management or VPN name.

Field ID	Field output	Default position
traceroute_output	traceroute Output	1

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Allowed

Return codes

- NO_IP_INTERFACE_MATCHES_CRITERIA
  No IP Interface or VLAN ID matches the defined criteria. Check if the IP Address and/or VLAN ID are correct.

- MORE_THAN_ONE_IP_INTERFACE_MATCHES
  More than one IP Interface matches the defined criteria.

Testing the traceroute to a remote IPv6 address

Use the **ipinterface_run_traceroute6** command to test connectivity to a remote IPv6 address using the ICMP trace-route mechanism.

```
ipinterface_run_traceroute6 localipaddress6=IPv6address remote6=remoteHost
```

Parameters

Name	Description	Mandatory
localipaddress6	IPv6 address of the IP interface for which the traceroute6 command is run.	Y
remote6	IPv6 address or DNS for the traceroute test.	Y

This command runs a route trace to the specified remote host through the specified IP interface. The IP address must be one of the IP addresses defined for iSCSI IP interfaces or the Management or VPN name.

Field ID	Field output	Default position
traceroute_output	traceroute Output	1

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Allowed

Return codes

- NO_IP_INTERFACE_MATCHES_CRITERIA
  No IP Interface or VLAN ID matches the defined criteria. Check if the IP Address and/or VLAN ID are correct.
- MORE_THAN_ONE_IP_INTERFACE_MATCHES
  More than one IP Interface matches the defined criteria.

Updating an IP interface

Use the **ipinterface_update** command to update the configuration of an IP interface.

```
ipinterface_update ipinterface=IPInterfaceName [ address=Address ] [ netmask=NetworkMask ]
 [ gateway=DefaultGateway ] [ address6=IPv6address ] [ gateway6=DefaultIPv6Gateway ]
 [ access_group=IPAccessGroupName ] [ mtu=MTU ] [ vlan_id=VlanID ]
```

Parameters

Name	Type	Description	Mandatory	Default
ipinterface	Object name	The name of the IP interface to be updated.	Y	N/A
address	N/A	IP address of the interface or a list of addresses for the Management and VPN interfaces.	N	Leaves the address unchanged.
netmask	N/A	Network mask of the interface.	N	Leaves the network mask unchanged.
gateway	N/A	IP address of the default gateway for this interface.	N	Leaves unchanged.
address6	N/A	IPv6 address of the interface or a list of addresses for the Management and VPN interfaces.	N	Leaves the address unchanged.
gateway6	N/A	IPv6 address of the default gateway for this interface.	N	Leaves unchanged.
mtu	Integer	Maximum Transmission Unit: The packet size that is supported by the connecting Ethernet switch. Valid values 1500 - 9000.	N	Keep unchanged.
access_group	Object name	The name of the IP access group used for IP filtering.	N	Keep unchanged.
vlan_id	Integer	The virtual LAN ID. Valid values 1-4094.	N	0

This command updates the configuration of an existing IP interface.

Fields that are not specified do not change their values.

The name of the interface may either be one of the previously defined IP interfaces for iSCSI, or Management for the management IP interface, or VPN for the VPN interface.

Management ports are dedicated for CLI and GUI communications, as well as for outgoing SNMP and SMTP connections. For management interfaces, the user must specify three IP addresses (equal to the number of potential managers, minus the number of management ports).

For VPN interfaces, the user must specify two IP addresses (equal to the number of VPN ports). All VPN addresses must reside on the same subnet.

The mtu parameter can only be used if VLAN is not enabled and this IP interface is the only one configured on the Ethernet port.

Example:

```
ipinterface_update ipinterface=management
```

Output:

```
Command completed successfully
```

Attention: Be careful when changing Ethernet port MTU. If management connectivity is on, or a VPN is being used, this might break existing application connectivity, unless the adjacent switch is properly configured. This would be the result of the switch using a fixed MTU size lower than the new size being defined on the system. For example, if the system sends 9000 byte packets while the switch can only receive packets up to 1500 bytes, the switch drops the extra packets.

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Allowed

Return codes

- IPINTERFACE_DOES_NOT_EXIST
  This IP interface name does not exist.
- IP_ADDRESS_ALREADY_USED_IN_ANOTHER_INTERFACE
  The IP address is already assigned to another interface.
- IPADDRESS_AND_GATEWAY_ARE_NOT_ON_SAME_SUBNET
  The IP address specified for the default gateway is not in the subnet of the IP interface.
- IPINTERFACE_MANAGEMENT_DIFFERENT_SUBNET
  All IP addresses management modules must be in the same subnet.
- IPINTERFACE_MANAGEMENT_MISSING_IPS
  The number of IP addresses specified is smaller than the number of management modules.
- IPINTERFACE_MANAGEMENT_TOO_MANY_IPS
  The number of IP addresses specified is larger than the number of management modules.
- MTU_TOO_LARGE
  The specified MTU value is too large.
- ILLEGAL_IPADDRESS
  An illegal IP address was entered.
- DUPLICATE_IPADDRESSES
  Duplicate IP addresses were specified.
- ILLEGAL_GATEWAY_IPADDRESS
  An illegal IP address was specified for the default gateway.
- ILLEGAL_IPV6ADDRESS
  An illegal IPv6 address was entered.
- DUPLICATE_IPV6ADDRESSES
  Duplicate IPv6 addresses were specified.
- ILLEGAL_GATEWAY_IPV6_ADDRESS
  An illegal IPv6 address was specified for the default gateway.
- IPV6ADDRESS_AND_GATEWAY_ARE_NOT_ON_SAME_SUBNET
  The IPv6 address specified for the default gateway is not in the subnet of the IP interface.
- IPV6_ADDRESS_ALREADY_USED_IN_ANOTHER_INTERFACE
  The IPv6 address is already assigned to another interface.

- IPINTERFACE_MANAGEMENT_MISSING_IPV6S
  The number of IPv6 addresses specified is smaller than the number of management modules.
- IPINTERFACE_MANAGEMENT_TOO_MANY_IPV6S
  The number of IPv6 addresses specified is larger than the number of management modules.
- IPINTERFACE_MANAGEMENT_DIFFERENT_IPV6_SUBNET
  All IPv6 addresses management modules must be in the same subnet.
- IP_ACCESS_GROUP_DOES_NOT_EXIST
  An IP access group with the specified name does not exist.
- IP_ACCESS_INVALID_INTERFACE_TYPE
  IP filtering is applied to an invalid interface (should be management or VPN).
- IPINTERFACE_VLAN_UPDATE_LIMITED_TO_MANAGEMENT_OR_VPN
  The VLAN ID can be updated only for management or VPN IP interfaces.
- PARAMETER_CANNOT_BE_UPDATED_ON_IPINTERFACE_WITH_VLAN
  The MTU, speed, module or port parameter can be used only if VLAN is not configured and this is the only IP interface on the physical port.
- ETHERNET_PORT_DOES_NOT_EXIST
  The Ethernet port does not exist.

Updating VLAN priority code points

Use the **ipinterface_vlan_update** command to update the priority code points of all IP interfaces which have a VLAN ID defined.

```
ipinterface_vlan_update vlan_pcp=VlanPcp
```

Parameters

Name	Type	Description	Mandatory
vlan_pcp	Integer	The VLAN priority code point (PCP). Valid values 0-7.	Y

This command updates the priority code points for all IP interfaces which have a VLAN ID defined.

Example:

```
ipinterface_vlan_update vlan_pcp=2
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Allowed

Defining a new IPSec connection

Use the **ipsec_connection_add** command to add a new IPSec connection.

```
ipsec_connection_add ipsec_connection=ConnectionName left=IPInterfaceName [ right_ip=RightIpAddress ]
 < passkey=PassKey | certificate=PemCertificate >
```

Parameters

Name	Type	Description	Mandatory	Default
ipsec_connection	N/A	The name of the IPSec connection to be added.	Y	N/A
left	Object name	The name of the IP interface to be used as the left side: `management` or `VPN`.	Y	N/A
right_ip	N/A	IP address of the right side.	N	Any
passkey	N/A	Secret password.	N	N/A
certificate	N/A	The public certificate or certificate chain of the new IPSec connection (see below for details).	N	N/A

This command defines a new IPSec connection between an IP interface and the right side.

IP interface can be either **management** or **VPN**. If specified:

- The address of the right side is IPv4 or IPv6; otherwise the right side can be any.
- The secret password must be shared between the left and the right sides.
- The certificate must contain a public key of the right side.

The certificate parameter:

The value of the certificate parameter is the content of a PEM file with asterisks instead of newlines. Chained certificates are supported. The total maximal length of a PEM file holding chained certificates (leaf first, root last) is 15360 characters (including the asterisk characters). In Windows, you can drag-and-drop a PEM file from the Windows Explorer to the appropriate location in the XCLI session window; the content will be added automatically.

Example:

```
ldap_update_server fqdn=ldap.example.com address=192.0.2.1
        remove_certificate=yes
```

Example:

```
ipsec_connection_add ipsec_connection=MySec left=management passkey="MyPass123"
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Allowed

Return codes

- IPSEC_CONNECTION_EXISTS
  The IPSec connection already exists.
- IPSEC_CONNECTION_BETWEEN_ENDPOINTS_EXISTS
  A connection between these endpoints already exists.
- LEFT_INTERFACE_NOT_FOUND
  The specified left side interface was not found.
- MAX_IPSEC_CONNECTIONS_REACHED
  The maximum allowed number of configured IPSec connections is already reached.
- IPSEC_UNSUPPORTED_FOR_ISCSI
  IPSec is not supported for iSCSI ports.
- SSL_CERTIFICATE_CHAIN_EMPTY
  No certificates were found in the input.
- SSL_CERTIFICATE_HAS_EXPIRED
  The SSL certificate has expired.
- SSL_CERTIFICATE_INVALID_FORMAT
  The SSL certificate format is invalid or corrupted.
- SSL_CERTIFICATE_ISSUER_NOT_FOUND
  The SSL certificate issuer was not found in the certificate chain.
- SSL_CERTIFICATE_NOT_YET_VALID
  The SSL certificate is not yet valid.
- SSL_CERTIFICATE_VERIFICATION_FAILED
  The SSL certificate chain verification failed.
- SSL_CERTIFICATE_VERIFICATION_INTERNAL_ERROR
  The SSL certificate verification has failed because of an internal system error.

Updating an existing IPSec connection

Use the **ipsec_connection_update** command to update an existing IPSec connection.

```
ipsec_connection_update ipsec_connection=ConnectionName [ left=IPInterfaceName ] [ right_ip=RightIpAddress ]
 [ passkey=PassKey | certificate=PemCertificate ]
```

Parameters

Name	Type	Description	Mandatory	Default
ipsec_connection	Object name	The name of the IPSec connection to be updated.	Y	N/A
left	Object name	The name of the IP interface to be used as left side: **management** or **VPN**.	N	None
right_ip	N/A	The IP address of the right side.	N	None
passkey	N/A	Pre-shared key.	N	None
certificate	N/A	The public certificate or certificate chain of the IPSec connection to be updated (see below for details).	N	None

This command updates an existing IPSec connection between an IP interface and the right side.

IP interface can be either **management** or **VPN**. If specified:

- The address of the right side is IPv4 or IPv6; otherwise the right side can be any.
- The pre-shared key must be shared between the left and the right sides.
- The certificate must contain a public key of the right side.

The certificate parameter:

The value of the certificate parameter is the content of a PEM file with asterisks instead of newlines. Chained certificates are supported. The total maximal length of a PEM file holding chained certificates (leaf first, root last) is 15360 characters (including the asterisk characters). In Windows, you can drag-and-drop a PEM file from the Windows Explorer to the appropriate location in the XCLI session window; the content will be added automatically.

Example:

```
ipsec_connection_update ipsec_connection=MySec passkey="MyNewPass!@#"
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Allowed

Return codes

- IPSEC_CONNECTION_DOES_NOT_EXIST
  The specified IPSec connection does not exist.
- IPSEC_CONNECTION_EXISTS
  The IPSec connection already exists.
- LEFT_INTERFACE_NOT_FOUND
  The specified left side interface was not found.
- IPSEC_UNSUPPORTED_FOR_ISCSI
  IPSec is not supported for iSCSI ports.
- SSL_CERTIFICATE_CHAIN_EMPTY
  No certificates were found in the input.
- SSL_CERTIFICATE_HAS_EXPIRED
  The SSL certificate has expired.
- SSL_CERTIFICATE_INVALID_FORMAT
  The SSL certificate format is invalid or corrupted.
- SSL_CERTIFICATE_ISSUER_NOT_FOUND
  The SSL certificate issuer was not found in the certificate chain.
- SSL_CERTIFICATE_NOT_YET_VALID
  The SSL certificate is not yet valid.
- SSL_CERTIFICATE_VERIFICATION_FAILED
  The SSL certificate chain verification failed.
- SSL_CERTIFICATE_VERIFICATION_INTERNAL_ERROR
  The SSL certificate verification has failed because of an internal system error.

Removing an existing IPSec connection

Use the **ipsec_connection_remove** command to remove an existing IPSec connection.

```
ipsec_connection_remove ipsec_connection=ConnectionName
```

Parameters

Name	Type	Description	Mandatory
ipsec_connection	Object name	The name of the IPSec connection to be updated.	Y

Example:

```
ipsec_connection_remove ipsec_connection=connect1
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Allowed

Return codes

- IPSEC_CONNECTION_DOES_NOT_EXIST
  The specified IPSec connection does not exist.

Listing IPSec connections

Use the **ipsec_connection_list** command to list all or specific IPSec connections.

```
ipsec_connection_list [ ipsec_connection=ConnectionName ]
```

Parameters

Name	Type	Description	Mandatory	Default
ipsec_connection	Object name	The IPSec connection(s) to be listed.	N	All IPsec connections

Field ID	Field output	Default position
name	IPSec Connection	1
type	Type	2
left	Left Interface	3
right_ip	Right Address	4

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing IPSec tunnels

Use the **ipsec_list_tunnels** command to list IPSec tunnels.

```
ipsec_list_tunnels [ ipsec_connection=ConnectionName ] [ left=IPInterfaceName ] [ left_ip=InterfaceIpAddress ]
 [ right_ip=RightIpAddress ] [ module=ComponentId ]
```

Parameters

Name	Type	Description	Mandatory	Default
ipsec_connection	Object name	Lists all IPSec tunnels of this IPSec connection.	N	IPSec tunnels of all IPsec connections
left	Object name	Lists all IPSec tunnels from this interface.	N	IPsec tunnels from any interface
left_ip	N/A	Lists all IPSec tunnels from this left IP.	N	IPsec tunnels from any left IP
right_ip	N/A	Lists all IPSec tunnels from this right IP.	N	IPsec tunnels to any right IP
module	N/A	Limits the listing to a specific module.	N	All modules

Field ID	Field output	Default position
name	IPSec Connection	1
type	Type	2
status	Status	3
left	Left Interface	4
left_ip	Left Address	5
right_ip	Right Address	6
module	Module	7

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Connecting to a support center

Use the **support_center_connect** command to connect to a support center.

```
support_center_connect [ < timeout=Timeout [ idle_timeout=IdleTimeout ] > | always_on=<yes|no> ]
 [ module=ModuleNumber ] [ password=Password ]
```

Parameters

Name	Type	Description	Mandatory	Default
timeout	N/A	Specifies the duration of the session. After the duration elapses, the session will be disconnected. Time is specified in hh:mm format.	N	none
idle_timeout	N/A	Specifies the idle time for the session after which it will be disconnected. Time is specified in hh:mm format.	N	[timeout]
module	N/A	The module from which the connection to the support center should be initiated	N	[the module that handled the CLI request]
password	String	A password set by the customer, that needs to be submitted by support services, in order to start a remote support session Format: string, must be 6-12 alpha-numeric characters, and is case-insensitive.	N	none
always_on	Boolean	Enables a constant connection to the support center (rather than an on-demand connection).	N	none

If the support center is not defined, the command will fail.

To control the duration of the session, use the parameters timeout and idle_disconnect.

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Allowed

Return codes

- NO_SUPPORT_CENTERS_ARE_DEFINED
  No support centers are defined.

- IDLE_TIMEOUT_MUST_BE_LOWER_THAN_TIMEOUT
  The idle timeout must be smaller than the regular timeout.
- MODULE_HAS_NO_SUPPORT_CENTER_PORT
  The specified module does not have a port from which the support center can connect.
- NO_MODULE_WITH_SUPPORT_CENTER_PORT
  No module has a port from which the support center can connect.
- REMOTE_SUPPORT_CLIENT_ALREADY_RUNNING
  The Remote Support Client is already running.
- REMOTE_SUPPORT_CLIENT_AUTOMATICALLY_CONNECT_IS_RUNNING
  The Remote Support Client is running in automatically connect mode.

  Troubleshooting: Run support_center_disconnect to stop it.

Defining a support center

Use the **support_center_define** command to define a support center.

```
support_center_define support_center=SupportCenterName address=Address [ port=port ] [ priority=priority ]
```

Parameters

Name	Type	Description	Mandatory	Default
support_center	Object name	The name of the support center server	Y	N/A
address	N/A	The IP address of the support center server	Y	N/A
port	Positive integer	The TCP port to connect to on the support center	N	22
priority	N/A	The priority of the support center (support centers with a higher priority will be connected first)	N	0

Example:

```
support_center_define support_center=somewhere address=192.0.2.1
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Allowed

Return codes

- MAX_SUPPORT_CENTERS_DEFINED
  The maximum number of defined support centers is already reached.
- SUPPORT_CENTER_ALREADY_EXISTS
  The support center with the indicated name already exists.

  Troubleshooting: Enter a different name.

Deleting a support center

Use the **support_center_delete** command to delete a support center.

```
support_center_delete support_center=SupportCenterName
```

Parameters

Name	Type	Description	Mandatory
support_center	Object name	The name of the support center to delete.	Y

Sessions that belong to this support center are disconnected, even if they are open at the time of deletion.

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Allowed

Warnings

- ARE_YOU_SURE_TO_DELETE_THE_SUPPORT_CENTER
  Are you sure you want to delete the support center?.

Return codes

- SUPPORT_CENTER_NOT_DEFINED
  The support center is not defined.
- CANNOT_DELETE_WHILE_SUPPORT_CENTER_IS_RUNNING
  The support center is running. Disconnect it before deleting.

Disconnecting from a support center

Use the **support_center_disconnect** command to disconnect the storage system from a support center.

```
support_center_disconnect
```

Example:

```
support_center_disconnect
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Allowed

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_DISCONNECT_BUSY_REMOTE_SUPPORT
  Are you sure you want to disconnect the busy remote support connection?

Return codes

- REMOTE_SUPPORT_CLIENT_NOT_RUNNING
  The Remote Support Client is not running.

Listing support centers

Use the **support_center_list** command to list support centers.

```
support_center_list
```

This command displays the following information about all defined support centers:

- Name
- IP Address
- Port

- Priority

Field ID	Field output	Default position
name	Name	1
address	Address	2
port	Port	3
priority	Priority	4

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing statuses of all support centers

Use the **support_center_status** command to list information about the status of the defined support centers.

```
support_center_status
```

Example:

```
support_center_status
```

Output:

```
State            Connected sessions   Timeout (min)   Module   Connected since
---------------- -------------------- --------------- -------- -----------------
no connection    0                    no timeout
Cont.:
Destination    Auto Connect Active   Always On
-------------- --------------------- -----------
               no                    no
```

Field ID	Field output	Default position
state	State	1
connected_support_sessions	Connected sessions	2
minutes_to_timeout	Timeout (min)	3
running_from_module	Module	4
start_time	Connected since	5
destination	Destination	6
automatically_connect_mode	Auto Connect Active	7
stop_automatically_connect	Stop support center automatically connect	N/A
always_on	Always On	8

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Allowed

Configuring the support center connection to enable automatic connect on restart

Use the **support_center_config** command to configure the connection to a support center automatically.

```
support_center_config automatically_connect=<yes|no> [ connect_through_module1=module ]
 [ connect_through_module2=module ] [ connect_through_module3=module ] [ password=Password ]
```

Parameters

Name	Type	Description	Mandatory	Default
automatically_connect	Boolean	Enables the automatic connection to the support center.	Y	N/A
connect_through_module1	N/A	The first module from which to establish a connection to the support center automatically.	N	Module with first management port
connect_through_module2	N/A	The second module from which to establish a connection to the support center automatically.	N	Module with second management port
connect_through_module3	N/A	The third module from which to establish a connection to the support center automatically.	N	Module with third management port
password	String	A password set by the customer, that needs to be submitted by support services, in order to start a remote support session Password format: case-insensitive string of 6-12 alphanumeric characters.	N	none

Example:

```
support_center_config automatically_connect=yes
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Allowed

Return codes

- MODULE_HAS_NO_SUPPORT_CENTER_PORT
  The specified module does not have a port from which the support center can connect.
- NO_MODULE_WITH_SUPPORT_CENTER_PORT
  No module has a port from which the support center can connect.

Listing the configuration of support center automatically connect

Use the **support_center_config_list** command to display the configuration of the automatic connection to a support center.

```
support_center_config_list
```

Example:

```
support_center_config_list
```

Output:

```
Enable Auto Conn   First Module   Second Module   Third Module
------------------ -------------- --------------- --------------
yes                1              2               -1
```

Field ID	Field output	Default position
enable_auto_conn	Enable Auto Conn	1
module1_id	First Module	2
module2_id	Second Module	3
module3_id	Third Module	4
automatically_connect_reason	Auto Conn Reason	N/A

Access control

User Category	Permission
Storage administrator	Allowed

User Category	Permission
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Creating a new IP access group

Use the **ip_access_group_create** command to create a new IP access group.

```
ip_access_group_create access_group=IPAccessGroupName
```

Parameters

Name	Type	Description	Mandatory
access_group	Object name	The name of the IP access group to be created.	Y

The group may contain up to 20 addresses and can be used to limit network access to a management/VPN interface.

Example:

```
ip_access_group_create access_group=IPAccessGroup1
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- IP_ACCESS_GROUP_ALREADY_EXISTS
  An IP access group with the specified name already exists.
- IP_ACCESS_MAXIMUM_NUMBER_OF_GROUPS_IS_REACHED
  The maximum number of IP access groups is already reached.

Removing an address from an IP access group

Use the **ip_access_group_remove_address** command to delete the IP address of an access group.

```
ip_access_group_remove_address access_group=IPAccessGroupName address=Address
```

Parameters

Name	Type	Description	Mandatory
access_group	Object name	The name of the IP access group.	Y
address	N/A	The address that should be deleted from the IP access group.	Y

As a prerequisite for completing this command, the IP address must be defined for the group.

Example:

```
ip_access_group_remove_address access_group=IPAccessGroup1 address=192.0.2.1
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- IP_ACCESS_GROUP_DOES_NOT_EXIST
  An IP access group with the specified name does not exist.
- IP_ACCESS_ADDRESS_IS_NOT_VALID
  The given address is not valid.
- IP_ACCESS_ADDRESS_IS_NOT_IN_GROUP
  The specified address is not in the group.

Adding a new address to an IP access group

Use the **ip_access_group_add_address** command to add a new IP to an access group.

```
ip_access_group_add_address access_group=IPAccessGroupName address=Address [ netmask=NetworkMask ]
```

Parameters

Name	Type	Description	Mandatory	Default
access_group	Object name	The name of an IP access group.	Y	N/A
address	N/A	A valid IP4 address or FQDN to be added to the IP access group.	Y	N/A
netmask	N/A	The network mask for a network address range.	N	Single IP address range (255.255.255.255).

The address can be an IP4 address with or without a netmask, or a valid host name (FQDN).

Example:

```
ip_access_group_add_address access_group=IPAccessGroup1 address=192.0.2.1
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- IP_ACCESS_GROUP_DOES_NOT_EXIST
  An IP access group with the specified name does not exist.
- IP_ACCESS_REMOTE_RESOLVE_ADDRESS_CALL_HAS_FAILED
  The remote call to resolve an address has failed.
- IP_ACCESS_MAXIMUM_NUMBER_OF_ADDRESSES_IN_GROUP_IS_REACHED
  The maximum number of addresses in the IP access group is already reached.

Deleting an existing IP access group

Use the **ip_access_group_delete** command to delete an IP access group.

```
ip_access_group_delete access_group=IPAccessGroupName
```

Parameters

Name	Type	Description	Mandatory
access_group	Object name	The name of the IP access group to be deleted.	Y

Example:

```
ip_access_group_delete access_group=DBGroupNew
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- IP_ACCESS_GROUP_DOES_NOT_EXIST
  An IP access group with the specified name does not exist.
- IP_ACCESS_GROUP_IN_USE
  The group is used for IP filtering.

Renaming an existing IP access group

Use the **ip_access_group_rename** command to rename an existing IP access group.

```
ip_access_group_rename access_group=IPAccessGroupName new_name=Name
```

Parameters

Name	Type	Description	Mandatory
access_group	Object name	Name of the IP access group to be renamed.	Y
new_name	Object name	A new name of the IP access group.	Y

Example:

```
ip_access_group_rename access_group=DBGroup new_name=DBGroupNew
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- IP_ACCESS_GROUP_DOES_NOT_EXIST
  An IP access group with the specified name does not exist.
- IP_ACCESS_GROUP_ALREADY_EXISTS
  An IP access group with the specified name already exists.

Listing IP access groups

Use the **ip_access_group_list** command to list IP access groups.

```
ip_access_group_list
```

Field ID	Field output	Default position
name	Group Name	1
addresses.0	Address 1	N/A
addresses.1	Address 2	N/A
addresses.2	Address 3	N/A
addresses.3	Address 4	N/A
addresses.4	Address 5	N/A
addresses.5	Address 6	N/A
addresses.6	Address 7	N/A
addresses.7	Address 8	N/A
addresses.8	Address 9	N/A
addresses.9	Address 10	N/A
addresses.10	Address 11	N/A
addresses.11	Address 12	N/A
addresses.12	Address 13	N/A
addresses.13	Address 14	N/A
addresses.14	Address 15	N/A
addresses.15	Address 16	N/A
addresses.16	Address 17	N/A
addresses.17	Address 18	N/A
addresses.18	Address 19	N/A
addresses.19	Address 20	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Allowed
Technicians	Disallowed

Listing IP access group addresses

Use the **ip_access_group_address_list** command to list IP access group addresses.

```
ip_access_group_address_list
```

This command lists IP access groups and address lists for these groups.

Example:

```
ip_access_group_address_list
```

Output:

```
Group Name       Address
---------------- ----------------
DBGroup          192.0.2.1
IPAccessGroup1   198.51.100.2
```

Field ID	Field output	Default position
access_group	Group Name	1
address	Address	2

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Allowed

User Category	Permission
Technicians	Disallowed

PKI configuration commands

This section describes the command-line interface (CLI) for PKI configuration.

- **Listing PKI items**
  Use the **pki_list** command to list PKI items.
- **Generating a certificate signing request**
  Use the **pki_generate_csr** command to generate a certificate signing request.
- **Generating a private key and CSR**
  Use the **pki_generate_private_key_and_csr** command to generate a private key and CSR.
- **Deleting a PKI content**
  Use the **pki_remove** command to delete the PKI content.
- **Changing a PKI symbolic name**
  Use the **pki_rename** command to change a PKI symbolic name.
- **Importing a signed certificate**
  Use the **pki_set_pem** command to import a signed certificate in PEM format.
- **Importing a PKCS#12 certificate**
  Use the **pki_set_pkcs12** command to import a PKCS#12 certificate.
- **Displaying the details of a signed certificate**
  Use the **pki_show_certificate** command to display the details of a signed certificate.
- **Updating a PKI certificate or services**
  Use the **pki_update** command to update a PKI certificate or services.

Listing PKI items

Use the **pki_list** command to list PKI items.

```
pki_list
```

The storage system allows you to install certificates generated by your own certificate authority (CA) for the different services that use digital certificates (SSL authentication, IPSec, and so on). When you install a certificate, it is associated with a name that you provide, which is used for managing it.

Certificates can be installed in one of two ways, depending on your site PKI policy:

- System generated: This method does not expose the system private key
  - The system generates a public-private keypair
  - The public key is exported in a certificate signing request (CSR) file using the pki_generate_private_key_and_csr command.
  - CA generated: The CA signs this file, returning a .PEM file that is then imported into the storage system using the pki_set_pem command.
- The CA generates both the key pair and associated certificate. Both are provided in a password-protected PKCS#12 file.
  - This file is imported into the system using the pki_set_pkcs12 command.

The pki_list command lists the following information:

- Name
- Fingerprint
- Has signed certificate
- Services

Field ID	Field output	Default position
name	Name	1
fingerprint	Fingerprint	2
authenticated	Has signed certificate	3
services	Services	4

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Generating a certificate signing request

Use the **pki_generate_csr** command to generate a certificate signing request.

```
pki_generate_csr name=Name subject=Subject
```

Parameters

Name	Type	Description	Mandatory
name	String	The certificate's symbolic name.	Y
subject	N/A	The subject name for the generated certificate request. The argument must be formatted as /type0=value0/type1=value1/type2=... .	Y

Example:

```
pki_generate_csr name subject
```

Field ID	Field output	Default position
csr	CSR	1

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- FAILED_CREATING_CERTIFICATE_SIGNING_REQUEST
  Failed to generate a certificate signing request.

  Troubleshooting: Generate a certificate signing request with a correct subject (for example, '/C=US/CN=IBM').

- CERTIFICATE_NAME_DOES_NOT_EXIST
  A certificate with the indicated name was not found.

  Troubleshooting: Enter a different name.

Generating a private key and CSR

Use the **pki_generate_private_key_and_csr** command to generate a private key and CSR.

```
pki_generate_private_key_and_csr name=Name subject=Subject [ bits=Bits ]
```

Parameters

Name	Type	Description	Mandatory	Default
bits	Integer	The private key size in bits. It can be between 1024 to 4096.	N	2048
name	String	The certificate's symbolic name.	Y	N/A
subject	N/A	The subject name for the generated certificate request. The argument must be formatted as /type0=value0/type1=value1/type2=... .	Y	N/A

Example:

```
pki_generate_private_key_and_csr name="my_cert"
subject="/C=US/CN=IBM" bits=1024
```

Field ID	Field output	Default position
csr	CSR	1

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- FAILED_CREATING_CERTIFICATE_SIGNING_REQUEST
  Failed to generate a certificate signing request.

  Troubleshooting: Generate a certificate signing request with a correct subject (for example, '/C=US/CN=IBM').
- FAILED_CREATING_PRIVATE_KEY
  Failed to create a private key.
- CERTIFICATE_NAME_ALREADY_EXIST
  A certificate with the indicated name already exists.

  Troubleshooting: Enter a different name.
- CERTIFICATE_CONTAINER_FULL
  Cannot add any more certificates, the maximum number is already reached.

  Troubleshooting: Delete a certificate.

Deleting a PKI content

Use the **pki_remove** command to delete the PKI content.

```
pki_remove name=Name
```

Parameters

Name	Type	Description	Mandatory
name	String	The certificate's symbolic name.	Y

Example:

```
pki_remove name="my_cert"
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_DELETE_CERTIFICATE
  Are you sure you want to delete the certificate?

Return codes

- CERTIFICATE_NAME_DOES_NOT_EXIST
  A certificate with the indicated name was not found.

  Troubleshooting: Enter a different name.
- DEFAULT_CERTIFICATE_CANNOT_BE_DELETED
  The default certificate cannot be deleted.

Changing a PKI symbolic name

Use the **pki_rename** command to change a PKI symbolic name.

```
pki_rename name=Name new_name=Name
```

Parameters

Name	Type	Description	Mandatory
name	String	The current symbolic name.	Y

Name	Type	Description	Mandatory
new_name	String	The new symbolic name.	Y

Example:

```
pki_rename name="current_name" new_name="my_new_name"
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- CERTIFICATE_NAME_ALREADY_EXIST
  A certificate with the indicated name already exists.

  Troubleshooting: Enter a different name.

- CERTIFICATE_NAME_DOES_NOT_EXIST
  A certificate with the indicated name was not found.

  Troubleshooting: Enter a different name.

Importing a signed certificate

Use the **pki_set_pem** command to import a signed certificate in PEM format.

```
pki_set_pem certificate=SignedCertificate  [ services=<xcli [ ,cim ] [ ,ipsec ] ... | ALL | NONE> ]
```

Parameters

Name	Description	Mandatory	Default
services	A comma-separated list of services that use this certificate.	N	none
certificate	The public certificate or certificate chain to be imported (see below for details).	Y	N/A

As a security precaution, use the **pki_show_security** command to view the certificate in plain text, and make sure that the certificate text under *Signature Algorithm* does not include the string *MD5*. This will help you avoid a "transcript collision" attack, that can force a hash-construction downgrade to MD5 and reduce expected security. For the vulnerability summary, see the National Vulnerability Database.

The certificate parameter

The value of the certificate parameter is the content of a PEM file with asterisks instead of newlines. Chained certificates are supported. The total maximal length of a PEM file holding chained certificates (leaf first, root last) is 15360 characters (including the asterisk characters). In Windows, you can drag-and-drop a PEM file from the Windows Explorer to the appropriate location in the CLI session window; the content will be added automatically.

Example:

```
pki_set_pem certificate=validCertificateChain
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- SERVICE_IS_USING_OTHER_CERTIFICATE
  Service '*services*' is using another certificate.

Troubleshooting: Edit the certificate previously used by the service.

- PRIVATE_KEY_ALREADY_HAS_OTHER_CERTIFICATE
  The private key matching this certificate already has another certificate.

  Troubleshooting: To replace the certificate, use the pki_update command.

- CERTIFICATE_KEY_WAS_NOT_FOUND
  Failed to set the certificate.

  Troubleshooting: Make sure the certificate parameters are correct.

- SSL_CERTIFICATE_CHAIN_EMPTY
  No certificates were found in the input.
- SSL_CERTIFICATE_HAS_EXPIRED
  The SSL certificate has expired.
- SSL_CERTIFICATE_INVALID_FORMAT
  The SSL certificate format is invalid or corrupted.
- SSL_CERTIFICATE_ISSUER_NOT_FOUND
  The SSL certificate issuer was not found in the certificate chain.
- SSL_CERTIFICATE_NOT_YET_VALID
  The SSL certificate is not yet valid.
- SSL_CERTIFICATE_VERIFICATION_FAILED
  The SSL certificate chain verification failed.
- SSL_CERTIFICATE_VERIFICATION_INTERNAL_ERROR
  The SSL certificate verification has failed because of an internal system error.

Importing a PKCS#12 certificate

Use the **pki_set_pkcs12** command to import a PKCS#12 certificate.

```
pki_set_pkcs12 name=Name password=Password certificate=Base64Data [ services=<xcli [ ,cim ] [ ,ipsec ]
 ... | ALL | NONE> ]
```

Parameters

Name	Type	Description	Mandatory	Default
services	N/A	A comma-separated list of services that use this certificate.	N	none
password	String	The PKCS#12 file password.	Y	N/A
name	String	The certificate's symbolic name.	Y	N/A
certificate	N/A	The PKCS#12 content in one-line base64 format. Such input can be created, for example, by a base64 utility: base64 -w0 myCert.pfx	Y	N/A

As a security precaution, use the **pki_show_security** command to view the certificate in plain text, and make sure that the certificate text under *Signature Algorithm* does not include the string *MD5*. This will help you avoid a "transcript collision" attack, that can force a hash-construction downgrade to MD5 and reduce expected security. For the vulnerability summary, see the National Vulnerability Database.

Note:
Quorum Witness communication might break in cases where replacing a certificate after HyperSwap relations have been created, if the default certificate has been overwritten while specifying the specific quorum service. This occurs when there are existing HyperSwap relations on the systems registered on the Quorum Witness and they are not deactived and before the certificate change. For more information, see the latest Quorum Witness user information within the FlashSystem A9000 or A9000R documentation set on IBM Documentation.

Example:

```
pki_set_pkcs12 name=myPki password=pkiPassword certificate=pkiCertificateBase64
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- SERVICE_IS_USING_OTHER_CERTIFICATE
  Service '*services*' is using another certificate.

  Troubleshooting: Edit the certificate previously used by the service.
- CANNOT_VALIDATE_PKCS12_FILE
  Failed validating PKCS#12 file.

  Troubleshooting: Make sure that the PKCS#12 file content is encoded to base64, and the password is correct.
- DEFAULT_CERTIFICATE_ALREADY_EXIST
  The default certificate already exists.

  Troubleshooting: Delete the default certificate or make it not default.
- CERTIFICATE_NAME_ALREADY_EXIST
  A certificate with the indicated name already exists.

  Troubleshooting: Enter a different name.
- BAD_BASE64_DATA
  Data cannot be decoded as base-64 data.
- FAILED_GETTING_PRIVATE_KEY_FINGERPRINT
  Failed to retrieve a private key fingerprint.
- FAILED_ENCRYPTING_PRIVATE_KEY
  Failed to encrypt a private key.
- CERTIFICATE_CONTAINER_FULL
  Cannot add any more certificates, the maximum number is already reached.

  Troubleshooting: Delete a certificate.

Displaying the details of a signed certificate

Use the **pki_show_certificate** command to display the details of a signed certificate.

```
pki_show_certificate name=Name
```

Parameters

Name	Type	Description	Mandatory
name	String	The certificate's symbolic name.	Y

As a security precaution, use this command to view the certificate in plain text, and make sure that the certificate text under *Signature Algorithm* does not include the string *MD5*. This will help you avoid a "transcript collision" attack, that can force a hash-construction downgrade to MD5 and reduce expected security. For the vulnerability summary, see the National Vulnerability Database.

Example:

```
pki_show_certificate name=ibm
```

Field ID	Field output	Default position
certificate	Certificate	1

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- FAILED_PARSING_CERTIFICATE
  Failed parsing the certificate.
- KEY_HAS_NO_CERTIFICATE
  The key has no signed certificate defined.

- CERTIFICATE_NAME_DOES_NOT_EXIST
  A certificate with the indicated name was not found.

  Troubleshooting: Enter a different name.

Updating a PKI certificate or services

Use the **pki_update** command to update a PKI certificate or services.

```
pki_update  name=Name [ services=<xcli [ ,cim ] [ ,ipsec ] ... | ALL | NONE> ]
 [ certificate=SignedCertificate ]
```

Parameters

Name	Type	Description	Mandatory	Default
services	N/A	Comma-separated list of services that need to use this certificate.	N	none
name	String	The certificate's symbolic name.	Y	N/A
certificate	N/A	If this parameter is defined, the certificate will be replaced.	N	none

Example:

```
pki_update name=cert services=xcli,cim
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- SERVICE_IS_USING_OTHER_CERTIFICATE
  Service '*services*' is using another certificate.

  Troubleshooting: Edit the certificate previously used by the service.
- NO_PKI_UPDATE_PARAMETERS_SPECIFIED
  No parameters were specified for the update.
- CERTIFICATE_DOES_NOT_MATCH_PRIVATE_KEY
  The certificate does not match the private key.

  Troubleshooting: Use another certificate.
- CANNOT_SET_SERVICES_BEFORE_SETTING_CERTIFICATE
  Cannot set services before setting the certificate.

  Troubleshooting: Set the certificate first.
- DEFAULT_CERTIFICATE_ALREADY_EXIST
  The default certificate already exists.

  Troubleshooting: Delete the default certificate or make it not default.
- CERTIFICATE_KEY_WAS_NOT_FOUND
  Failed to set the certificate.

  Troubleshooting: Make sure the certificate parameters are correct.
- CERTIFICATE_NAME_DOES_NOT_EXIST
  A certificate with the indicated name was not found.

  Troubleshooting: Enter a different name.
- SSL_CERTIFICATE_CHAIN_EMPTY
  No certificates were found in the input.
- SSL_CERTIFICATE_HAS_EXPIRED
  The SSL certificate has expired.
- SSL_CERTIFICATE_INVALID_FORMAT
  The SSL certificate format is invalid or corrupted.

- SSL_CERTIFICATE_ISSUER_NOT_FOUND
  The SSL certificate issuer was not found in the certificate chain.
- SSL_CERTIFICATE_NOT_YET_VALID
  The SSL certificate is not yet valid.
- SSL_CERTIFICATE_VERIFICATION_FAILED
  The SSL certificate chain verification failed.
- SSL_CERTIFICATE_VERIFICATION_INTERNAL_ERROR
  The SSL certificate verification has failed because of an internal system error.

InfiniBand commands

This section describes the command-line interface (CLI) for InfiniBand fabric management.

- **Listing the configured InfiniBand ports**
  Use the **ib_port_list** command to list the configured InfiniBand ports.
- **Listing data counters for the enabled InfiniBand switch ports**
  Use the **ib_port_counter_list** command to list data counters for the enabled InfiniBand switch ports.
- **Listing error counters for enabled InfiniBand switch ports**
  Use the **ib_port_error_list** command to list error counters for the enabled InfiniBand switch ports.
- **Listing the status of the enabled InfiniBand switch ports**
  Use the **ib_port_info_list** command to list the status of the enabled InfiniBand switch ports.
- **Listing data counters for the enabled InfiniBand HCA ports**
  Use the **ib_hca_counter_list** command to list data counters for InfiniBand HCA ports, enabled on modules and flash enclosures.
- **Listing error counters for the enabled InfiniBand HCA ports.**
  Use the **ib_hca_error_list** command to list error counters for the enabled InfiniBand HCA ports.
- **Listing the statuses of the enabled InfiniBand HCA ports**
  Use the **ib_hca_info_list** command to list the statuses of InfiniBand HCA ports, enabled on modules and flash enclosures.
- **Listing the configured InfiniBand switches**
  Use the **switch_list** command to list the configured InfiniBand switches.
- **Listing the configured InfiniBand switch management addresses**
  Use the **switch_mgmt_ip_list** command to list the configured InfiniBand switch management addresses.
- **Listing the configured InfiniBand switch firmware versions**
  Use the **switch_fw_list** command to list the configured InfiniBand switch firmware versions.
- **Listing the configured InfiniBand switch power values**
  Use the **switch_power_list** command to list the configured InfiniBand switch power values for PSUs and BBUs.
- **Listing the configured InfiniBand switch voltage values**
  Use the **switch_voltage_list** command to list the configured InfiniBand switch voltage values.
- **Listing the configured InfiniBand switch temperature values**
  Use the **switch_temp_list** command to list the configured InfiniBand switch temperature values.
- **Listing the configured InfiniBand switch fan parts**
  Use the **switch_fan_part_list** command to list the configured InfiniBand switch fan parts.
- **Listing the configured InfiniBand switch PSUs**
  Use the **switch_psu_list** command to list the configured InfiniBand switch PSUs.
- **Listing the configured InfiniBand switch BBUs**
  Use the **switch_bbu_list** command to list the configured InfiniBand switch BBUs.
- **Listing the configured InfiniBand switch fans**
  Use the **switch_fan_list** command to list the configured InfiniBand switch fans.

Listing the configured InfiniBand ports

Use the **ib_port_list** command to list the configured InfiniBand ports.

```
ib_port_list [ ib_port=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
ib_port	The InfiniBand port to be listed.	N	All IB ports

Example:

```
ib_port_list
```

Field ID	Field output	Default position
port	Port	1
component_id	Connected Component	2
status	Status	3
skip_miswire	Allow Any GUID	4
saved_info.peer_guid	GUID	5

Field ID	Field output	Default position
saved_info.last_state	State	6
saved_info.is_cm_ok	CM OK	N/A
saved_info.port_down_reason	Failure Reason	7
saved_info.last_state_change	Last State Change	N/A
saved_info.last_cm_check	Last CM Check	N/A
pending_ia_cmd	Component Operation	N/A
currently_functioning	Currently Functioning	N/A
requires_service	Requires Service	N/A
service_reason	Service Reason	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing data counters for the enabled InfiniBand switch ports

Use the **ib_port_counter_list** command to list data counters for the enabled InfiniBand switch ports.

```
ib_port_counter_list [ ib_port=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
ib_port	The InfiniBand switch port to be listed.	N	All IB switch ports

Example:

```
ib_port_counter_list
```

Output:

```
Port                    TX Data          RX Data          TX Pkt
----------------------- ---------------- ---------------- --------------
1:IB_Switch_Port:1:12   32110694059954   34132513631987   252139221702
1:IB_Switch_Port:1:13   10294316107140   13659311859037   72287616839

RX Pkt          XmtWait
--------------- -----------
251995529475    752259049
322306678848    41571936
```

Field ID	Field output	Default position
port	Port	1
XmtData	TX Data	2
RcvData	RX Data	3
XmtPkts	TX Pkt	4
RcvPkts	RX Pkt	5
XmtWait	XmtWait	6

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing error counters for enabled InfiniBand switch ports

Use the **ib_port_error_list** command to list error counters for the enabled InfiniBand switch ports.

```
ib_port_error_list [ ib_port=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
ib_port	The InfiniBand switch port to be listed.	N	All IB switch ports

Example:

```
ib_port_error_list
```

Output:

```
Port                    SymErr   LinkRec   LinkDown   RcvErr   RcvRPErr   RcvSRErr
----------------------- -------- --------- ---------- -------- ---------- ----------
1:IB_Switch_Port:1:12   0        0         134        0        0          31778
1:IB_Switch_Port:1:13   0        0         134        0        0          5514

XmtDisc   XmtCErr   RcvCErr   LinkIErr   ExcBOEre   VL15Dr
--------- --------- --------- ---------- ---------- --------
20237     0         0         0          0          0
35740     0         0         0          0          0
```

Field ID	Field output	Default position
port	Port	1
SymbolErrors	SymErr	2
LinkRecovers	LinkRec	3
LinkDowned	LinkDown	4
RcvErrors	RcvErr	5
RcvRemotePhysErrors	RcvRPErr	6
RcvSwRelayErrors	RcvSRErr	7
XmtDiscards	XmtDisc	8
XmtConstraintErrors	XmtCErr	9
RcvConstraintErrors	RcvCErr	10
LinkIntegrityErrors	LinkIErr	11
ExcBufOverrunErrors	ExcBOEre	12
VL15Dropped	VL15Dr	13

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing the status of the enabled InfiniBand switch ports

Use the **ib_port_info_list** command to list the status of the enabled InfiniBand switch ports.

```
ib_port_info_list [ ib_port=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
ib_port	The InfiniBand switch port to be listed.	N	All IB switch ports

Example:

```
ib_port_info_list
```

Output:

```
Port                    IB Log State   IB Phys State   Link Speed     Link Width
----------------------- -------------- --------------- -------------- -------------
1:IB_Switch_Port:1:12   ACTIVE         LINK UP         14.0625 Gbps   X4
1:IB_Switch_Port:1:13   ACTIVE         LINK UP         14.0625 Gbps   X4
```

Field ID	Field output	Default position
port	Port	1
log_state	IB Log State	2
phys_state	IB Phys State	3
link_speed	Link Speed	4
link_width	Link Width	5
link_width_sup	Link Width Sup	N/A
link_speed_sup	Link Speed Sup	N/A
link_speed_enabled	Link Speed Ena	N/A
link_width_enabled	Link Width Ena	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing data counters for the enabled InfiniBand HCA ports

Use the **ib_hca_counter_list** command to list data counters for InfiniBand HCA ports, enabled on modules and flash enclosures.

```
ib_hca_counter_list [ hca_port=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
hca_port	The InfiniBand HCA port to be listed.	N	All IB HCA ports

Example:

```
ib_hca_counter_list
```

Output:

```
Port                          TX Data         RX Data         TX Pkt
----------------------------- --------------- --------------- ------------
1:IB_FlashSystem_Port:4:1     0               0               0
1:IB_FlashSystem_Port:4:3     0               0               0
1:IB_FlashSystem_Port:4:5     0               0               0
1:IB_FlashSystem_Port:4:7     0               0               0
1:IB_Module_Port:12:1         254584779527    252260909954    2055238854
1:IB_Module_Port:12:2         0               0               0
1:IB_Module_Port:13:1         252395242864    254798454598    2061534883
1:IB_Module_Port:13:2         0               0               0
1:IB_Module_Port:8:1          254003578209    254027205845    2055494787
1:IB_Module_Port:8:2          0               0               0

RX Pkt        XmtWait
------------- -----------
0             0
0             0
0             0
0             0
2058771428    103686442
0             0
2060782961    94235849
0             0
2059021166    103031319
0             0
```

Field ID	Field output	Default position
port	Port	1
XmtData	TX Data	2
RcvData	RX Data	3
XmtPkts	TX Pkt	4
RcvPkts	RX Pkt	5
XmtWait	XmtWait	6

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing error counters for the enabled InfiniBand HCA ports.

Use the **ib_hca_error_list** command to list error counters for the enabled InfiniBand HCA ports.

```
ib_hca_error_list [ hca_port=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
hca_port	The InfiniBand HCA port to be listed.	N	All IB HCA ports

Example:

```
ib_hca_error_list
```

Field ID	Field output	Default position
port	Port	1
SymbolErrors	SymErr	2
LinkRecovers	LinkRec	3
LinkDowned	LinkDown	4
RcvErrors	RcvErr	5
RcvRemotePhysErrors	RcvRPErr	6
RcvSwRelayErrors	RcvSRErr	7
XmtDiscards	XmtDisc	8
XmtConstraintErrors	XmtCErr	9
RcvConstraintErrors	RcvCErr	10
LinkIntegrityErrors	LinkIErr	11
ExcBufOverrunErrors	ExcBOEre	12
VL15Dropped	VL15Dr	13

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing the statuses of the enabled InfiniBand HCA ports

Use the **ib_hca_info_list** command to list the statuses of InfiniBand HCA ports, enabled on modules and flash enclosures.

```
ib_hca_info_list [ hca_port=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
hca_port	The InfiniBand HCA port to be listed.	N	All IB HCA ports

Example:

```
ib_hca_info_list
```

Output:

```
Port                          IB Log State   IB Phys State
----------------------------- -------------- ---------------
1:IB_FlashSystem_Port:4:1     NOT SAMPLED    NOT SAMPLED
1:IB_FlashSystem_Port:4:3     NOT SAMPLED    NOT SAMPLED
1:IB_FlashSystem_Port:4:5     NOT SAMPLED    NOT SAMPLED
1:IB_FlashSystem_Port:4:7     NOT SAMPLED    NOT SAMPLED
1:IB_Module_Port:12:1         ACTIVE         LINK UP
1:IB_Module_Port:12:2         INIT           LINK UP
1:IB_Module_Port:13:1         ACTIVE         LINK UP
1:IB_Module_Port:13:2         INIT           LINK UP
1:IB_Module_Port:8:1          ACTIVE         LINK UP
1:IB_Module_Port:8:2          INIT           LINK UP
```

```
Link Speed     Link Width
-------------- --------------
NOT SAMPLED    NOT SAMPLED
NOT SAMPLED    NOT SAMPLED
NOT SAMPLED    NOT SAMPLED
NOT SAMPLED    NOT SAMPLED
14.0625 Gbps   X4
14.0625 Gbps   X4
14.0625 Gbps   X4
14.0625 Gbps   X4
14.0625 Gbps   X4
14.0625 Gbps   X4
```

Field ID	Field output	Default position
port	Port	1
log_state	IB Log State	2
phys_state	IB Phys State	3
link_speed	Link Speed	4
link_width	Link Width	5
link_width_sup	Link Width Sup	N/A
link_speed_sup	Link Speed Sup	N/A
link_speed_enabled	Link Speed Ena	N/A
link_width_enabled	Link Width Ena	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing the configured InfiniBand switches

Use the **switch_list** command to list the configured InfiniBand switches.

```
switch_list [ switch=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
switch	IB switch to list.	N	All IB switches

Example:

```
switch_list
```

Output:

```
Switch           Switch GUID         Status   MGMT OK   Ports OK   Power OK   BBU OK
---------------- ------------------- -------- --------- ---------- ---------- --------
1:IB_Switch:1    E41D2D03003C9900    OK       yes       yes        yes        yes
1:IB_Switch:2    E41D2D03003C9A80    OK       yes       yes        yes        yes
```

```
Cont.:
```

```
Fan OK   Temp OK   Volt OK   Boot Time             FW         Serial
-------- --------- --------- --------------------- ---------- --------------
yes      yes       yes       04/04/2016 10:51:18   3.5.0500   MT1523X09088
```

```
yes       yes       yes       27/03/2016 15:28:31   3.5.0500   MT1523X09091
```

Field ID	Field output	Default position
component_id	Switch	1
status	Status	3
sw_mgmt_status	MGMT Status	N/A
num_of_down_ports	Down Ports	N/A
mgmt_ok	MGMT OK	4
ports_ok	Ports OK	5
power_ok	Power OK	6
bbu_ok	BBU OK	7
fan_ok	Fan OK	8
temp_ok	Temp OK	9
volt_ok	Volt OK	10
fw	FW	12
mgmt_serial_number	Serial	13
mgmt_part_number	Part No	N/A
mgmt_asic_rev	ASIC Rev	N/A
mgmt_hw_rev	HW Rev	N/A
cpld_tor	CPLD Tor	N/A
cpld_port1	CPLD Port1	N/A
cpld_switch_brd	CPLD Switch Brd	N/A
chassis_serial_number	Chassis Serial	N/A
chassis_part_number	Chassis Part No	N/A
chassis_asic_rev	Chassis ASIC Rev	N/A
chassis_hw_rev	Chassis HW Rev	N/A
original_mgmt_serial_number	Original Serial	N/A
original_mgmt_part_number	Original Part No	N/A
original_mgmt_asic_rev	Original ASIC Rev	N/A
original_mgmt_hw_rev	Original HW Rev	N/A
original_chassis_serial_number	Original Chassis Serial	N/A
original_chassis_part_number	Original Chassis Part No	N/A
original_chassis_asic_rev	Original Chassis ASIC Rev	N/A
original_chassis_hw_rev	Original Chassis HW Rev	N/A
currently_functioning	Currently Functioning	N/A
mgmt_guid	Managememt GUID	N/A
requires_service	Requires Service	N/A
service_reason	Service Reason	N/A
fabric_switch_info.miss_count	Miss Count	N/A
fabric_switch_info.num_of_ports	Ports	N/A
fabric_switch_info.image_guid	FW GUID	N/A
fabric_switch_info.node_guid	Switch GUID	2
fabric_switch_info.dev_id	Device ID	N/A
fabric_switch_info.dev_rev	Device Revision	N/A
fabric_switch_info.vendor_id	Vendor ID	N/A
fabric_switch_info.name	Name	N/A
fabric_switch_info.mlx_dev_id	Ext Device ID	N/A
fabric_switch_info.mlx_hw_rev	Ext Device Revision	N/A
fabric_switch_info.boot_time	Boot Time	11
fabric_switch_info.uptime_seconds	Uptime	N/A
fabric_switch_info.fw_build_id	FW BUILD ID	N/A
fabric_switch_info.fw_rev	FW Ver	N/A
fabric_switch_info.fw_build_date	FW Build Date	N/A
fabric_switch_info.psid	PSID	N/A
used_power	Total Power Used	N/A
power_capacity	Total Power Capacity	N/A
power_available	Total Power Available	N/A
projected_max_used_power	Projected Max User Power	N/A
bbu_runtime	Battery Runtime	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed

User Category	Permission
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing the configured InfiniBand switch management addresses

Use the **switch_mgmt_ip_list** command to list the configured InfiniBand switch management addresses.

```
switch_mgmt_ip_list [ switch=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
switch	IB switch to list.	N	All IB switches

Example:

```
switch_mgmt_ip_list
```

Output:

```
Switch           Type     NAT IP         Router           Real IP          Status
---------------- -------- -------------- ---------------- ---------------- --------
1:IB_Switch:1    IPOIB    14.10.255.1    1:IB_Switch:1    14.10.255.1      OK
1:IB_Switch:1    MGMT1    14.10.254.1    1:Module:1       192.168.0.254    OK
1:IB_Switch:1    MGMT2    14.10.253.1    1:Module:4       192.168.1.254    OK
1:IB_Switch:1    SERIAL   14.10.10.3     1:Module:3       14.10.10.3       OK
1:IB_Switch:2    IPOIB    14.10.255.2    1:IB_Switch:2    14.10.255.2      OK
1:IB_Switch:2    MGMT1    14.10.254.2    1:Module:2       192.168.0.254    OK
1:IB_Switch:2    MGMT2    14.10.253.2    1:Module:3       192.168.1.254    OK
1:IB_Switch:2    SERIAL   14.10.10.4     1:Module:4       14.10.10.4       OK
```

Field ID	Field output	Default position
switch_id	Switch	1
type	Type	2
nat_ip	NAT IP	3
router	Router	4
real_ip	Real IP	5
status	Status	6

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing the configured InfiniBand switch firmware versions

Use the **switch_fw_list** command to list the configured InfiniBand switch firmware versions.

```
switch_fw_list [ switch=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
switch	IB switch to list.	N	All IB switches

Example:

```
switch_fw_list
```

Output:

```
Switch           Type                Version
---------------- ------------------- ----------
1:IB_Switch:1    ASIC                9.3.7170
1:IB_Switch:1    BBU1                703
1:IB_Switch:1    BBU2                703
1:IB_Switch:1    BIOS                4.6.5
1:IB_Switch:1    CPLD_PORT1          4
1:IB_Switch:1    CPLD_SWITCH_BRD     7
1:IB_Switch:1    CPLD_TOR            9
1:IB_Switch:1    MGMT                3.5.0500
1:IB_Switch:1    PSU1                404
1:IB_Switch:1    PSU2                404
1:IB_Switch:2    ASIC                9.3.7170
1:IB_Switch:2    BBU1                703
1:IB_Switch:2    BBU2                703
1:IB_Switch:2    BIOS                4.6.5
1:IB_Switch:2    CPLD_PORT1          4
1:IB_Switch:2    CPLD_SWITCH_BRD     7
1:IB_Switch:2    CPLD_TOR            9
1:IB_Switch:2    MGMT                3.5.0500
1:IB_Switch:2    PSU1                404
1:IB_Switch:2    PSU2                404
```

Field ID	Field output	Default position
switch_id	Switch	1
type	Type	2
version	Version	3
original_version	Original Version	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing the configured InfiniBand switch power values

Use the **switch_power_list** command to list the configured InfiniBand switch power values for PSUs and BBUs.

```
switch_power_list [ switch=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
switch	IB switch to list.	N	All IB switches

Example:

```
switch_power_list
```

Output:

```
Switch           Type   Power W   Voltage V   Current A   Capacity W   Feed   Status
---------------- ------ --------- ----------- ----------- ------------ ------ --------
1:IB_Switch:1    BBU1   N/A       N/A         N/A         330.00       NA     OK
1:IB_Switch:1    BBU2   N/A       N/A         N/A         330.00       NA     OK
1:IB_Switch:1    PS1    46.00     12.11       2.56        400.00       AC     OK
1:IB_Switch:1    PS2    44.00     12.19       2.75        400.00       AC     OK
1:IB_Switch:2    BBU1   N/A       N/A         N/A         330.00       NA     OK
1:IB_Switch:2    BBU2   N/A       N/A         N/A         330.00       NA     OK
1:IB_Switch:2    PS1    47.00     12.05       2.56        400.00       AC     OK
1:IB_Switch:2    PS2    44.00     12.14       2.81        400.00       AC     OK
```

Field ID	Field output	Default position
switch_id	Switch	1
type	Type	2
power	Power W	3
voltage	Voltage V	4
current	Current A	5
capacity	Capacity W	6

Field ID	Field output	Default position
feed	Feed	7
status	Status	8

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing the configured InfiniBand switch voltage values

Use the **switch_voltage_list** command to list the configured InfiniBand switch voltage values.

```
switch_voltage_list [ switch=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
switch	IB switch to list.	N	All IB switches

Example:

```
switch_voltage_list
```

Output:

```
Switch           Type             Expected   Actual   Status   High    Low
---------------- ---------------- ---------- -------- -------- ------- -------
1:IB_Switch:1    1.05V LAN        1.50       1.49     OK       1.72    1.27
1:IB_Switch:1    Asic 1.2V        1.20       1.20     OK       1.38    1.02
1:IB_Switch:1    Asic 1.8V        1.80       1.81     OK       2.06    1.53
1:IB_Switch:1    Asic 3.3V        3.30       3.31     OK       3.79    2.80
1:IB_Switch:1    BBU1             12.00      12.50    OK       13.80   10.19
1:IB_Switch:1    BBU2             12.00      12.50    OK       13.80   10.19
1:IB_Switch:1    CPU 0.9V         0.90       0.85     OK       1.03    0.77
1:IB_Switch:1    CPU 1.05V        1.05       1.03     OK       1.21    0.89
1:IB_Switch:1    CPU 1.8V         1.80       1.78     OK       2.06    1.53
1:IB_Switch:1    CPU/PCH 1.05V    1.05       1.00     OK       1.10    0.81
1:IB_Switch:1    DDR3 0.675V      0.68       0.66     OK       0.78    0.56
1:IB_Switch:1    DDR3 1.35V       1.35       1.34     OK       1.55    1.14
1:IB_Switch:1    PS1 vout 12V     12.00      12.11    OK       13.80   10.19
1:IB_Switch:1    PS2 vout 12V     12.00      12.19    OK       13.80   10.19
1:IB_Switch:1    SYS 3.3V         3.30       3.31     OK       3.79    2.80
1:IB_Switch:1    USB 5V           5.00       5.01     OK       5.75    4.25
1:IB_Switch:1    Vcore SX         0.95       0.96     OK       1.09    0.81
1:IB_Switch:2    1.05V LAN        1.50       1.52     OK       1.72    1.27
1:IB_Switch:2    Asic 1.2V        1.20       1.21     OK       1.38    1.02
1:IB_Switch:2    Asic 1.8V        1.80       1.81     OK       2.06    1.53
1:IB_Switch:2    Asic 3.3V        3.30       3.32     OK       3.79    2.80
1:IB_Switch:2    BBU1             12.00      12.50    OK       13.80   10.19
1:IB_Switch:2    BBU2             12.00      12.50    OK       13.80   10.19
1:IB_Switch:2    CPU 0.9V         0.90       0.86     OK       1.03    0.77
1:IB_Switch:2    CPU 1.05V        1.05       1.06     OK       1.21    0.89
1:IB_Switch:2    CPU 1.8V         1.80       1.83     OK       2.06    1.53
1:IB_Switch:2    CPU/PCH 1.05V    1.05       1.02     OK       1.10    0.81
1:IB_Switch:2    DDR3 0.675V      0.68       0.68     OK       0.78    0.56
1:IB_Switch:2    DDR3 1.35V       1.35       1.37     OK       1.55    1.14
1:IB_Switch:2    PS1 vout 12V     12.00      12.05    OK       13.80   10.19
1:IB_Switch:2    PS2 vout 12V     12.00      12.14    OK       13.80   10.19
1:IB_Switch:2    SYS 3.3V         3.30       3.41     OK       3.79    2.80
1:IB_Switch:2    USB 5V           5.00       5.16     OK       5.75    4.25
1:IB_Switch:2    Vcore SX         0.95       0.96     OK       1.09    0.81
```

Field ID	Field output	Default position
switch_id	Switch	1
type	Type	2
expected	Expected	3
actual	Actual	4
status	Status	5
high_margin	High	6
low_margin	Low	7

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing the configured InfiniBand switch temperature values

Use the **switch_temp_list** command to list the configured InfiniBand switch temperature values.

```
switch_temp_list [ switch=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
switch	IB switch to list.	N	All IB switches

Note:
The temperature values are indicated in Celsius.

Example:

```
switch_temp_list
```

Output:

```
Switch           Type          Actual   Alert    Critical   Status
---------------- ------------ -------- -------- ---------- --------
1:IB_Switch:1    BBU1          19.60    60.00    60.00      OK
1:IB_Switch:1    BBU2          19.50    60.00    60.00      OK
1:IB_Switch:1    MGMT_AMB      20.50    120.00   120.00     OK
1:IB_Switch:1    MGMT_CPU      25.00    120.00   120.00     OK
1:IB_Switch:1    MGMT_CPU1     22.00    120.00   120.00     OK
1:IB_Switch:1    MGMT_CPU2     25.00    120.00   120.00     OK
1:IB_Switch:1    MGMT_PORTS    22.00    120.00   120.00     OK
1:IB_Switch:1    MGMT_SX       28.00    105.00   110.00     OK
1:IB_Switch:1    PS1           24.00    120.00   120.00     OK
1:IB_Switch:1    PS2           24.00    120.00   120.00     OK
1:IB_Switch:2    BBU1          20.00    60.00    60.00      OK
1:IB_Switch:2    BBU2          19.89    60.00    60.00      OK
1:IB_Switch:2    MGMT_AMB      21.00    120.00   120.00     OK
1:IB_Switch:2    MGMT_CPU      26.00    120.00   120.00     OK
1:IB_Switch:2    MGMT_CPU1     25.00    120.00   120.00     OK
1:IB_Switch:2    MGMT_CPU2     19.00    120.00   120.00     OK
1:IB_Switch:2    MGMT_PORTS    22.50    120.00   120.00     OK
1:IB_Switch:2    MGMT_SX       28.00    105.00   110.00     OK
1:IB_Switch:2    PS1           24.00    120.00   120.00     OK
1:IB_Switch:2    PS2           25.00    120.00   120.00     OK
```

Field ID	Field output	Default position
switch_id	Switch	1
type	Type	2
actual	Actual	3
alert	Alert	4
critical	Critical	5
status	Status	6

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing the configured InfiniBand switch fan parts

Use the **switch_fan_part_list** command to list the configured InfiniBand switch fan parts.

```
switch_fan_part_list [ switch=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
switch	IB switch to list.	N	All IB switches

Example:

```
switch_fan_part_list
```

Output:

```
Switch           Type       Part No           Serial No       HW Rev   Speed      Status
---------------- ---------- ----------------- --------------- -------- ---------- --------
1:IB_Switch:1    FAN1-F1    MTEF-FANF-B       MT1523X09083    A2       10608.00   OK
1:IB_Switch:1    FAN1-F2    MTEF-FANF-B       MT1523X09083    A2       8998.00    OK
1:IB_Switch:1    FAN2-F1    MTEF-FANF-B       MT1523X09075    A2       10526.00   OK
1:IB_Switch:1    FAN2-F2    MTEF-FANF-B       MT1523X09075    A2       8939.00    OK
1:IB_Switch:1    FAN3-F1    MTEF-FANF-B       MT1523X09082    A2       10608.00   OK
1:IB_Switch:1    FAN3-F2    MTEF-FANF-B       MT1523X09082    A2       8998.00    OK
1:IB_Switch:1    FAN4-F1    MTEF-FANF-B       MT1523X09077    A2       10691.00   OK
1:IB_Switch:1    FAN4-F2    MTEF-FANF-B       MT1523X09077    A2       9242.00    OK
1:IB_Switch:1    PSU1       MTEF-PSF-AC-B     MT1523X09040    A5       14464.00   OK
1:IB_Switch:1    PSU2       MTEF-PSF-AC-B     MT1523X09044    A5       14592.00   OK
1:IB_Switch:2    FAN1-F1    MTEF-FANF-B       MT1523X09065    A2       10526.00   OK
1:IB_Switch:2    FAN1-F2    MTEF-FANF-B       MT1523X09065    A2       9118.00    OK
1:IB_Switch:2    FAN2-F1    MTEF-FANF-B       MT1523X09072    A2       10445.00   OK
1:IB_Switch:2    FAN2-F2    MTEF-FANF-B       MT1523X09072    A2       9118.00    OK
1:IB_Switch:2    FAN3-F1    MTEF-FANF-B       MT1523X09062    A2       10445.00   OK
1:IB_Switch:2    FAN3-F2    MTEF-FANF-B       MT1523X09062    A2       8998.00    OK
1:IB_Switch:2    FAN4-F1    MTEF-FANF-B       MT1523X09073    A2       10526.00   OK
1:IB_Switch:2    FAN4-F2    MTEF-FANF-B       MT1523X09073    A2       8998.00    OK
1:IB_Switch:2    PSU1       MTEF-PSF-AC-B     MT1523X09035    A5       14464.00   OK
1:IB_Switch:2    PSU2       MTEF-PSF-AC-B     MT1523X09033    A5       14464.00   OK
```

Field ID	Field output	Default position
switch_id	Switch	1
type	Type	2
part_no	Part No	3
serial_no	Serial No	4
hw_rev	HW Rev	5
speed	Speed	6
status	Status	7

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing the configured InfiniBand switch PSUs

Use the **switch_psu_list** command to list the configured InfiniBand switch PSUs.

```
switch_psu_list [ switch_psu=ComponentId ] [ switch=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
switch_psu	IB switch PSU to list.	N	All IB switch PSUs
switch	IB switch whose BBUs are to be listed.	N	All IB switches

Example:

```
switch_psu_list
```

Output:

```
Component ID          Status   Sensor Status   Power W   Voltage V
--------------------- -------- --------------- --------- -----------
1:IB_Switch_PSU:1:1   OK       OK              46.00     12.11
1:IB_Switch_PSU:1:2   OK       OK              44.00     12.19
1:IB_Switch_PSU:2:1   OK       OK              47.00     12.05
1:IB_Switch_PSU:2:2   OK       OK              44.00     12.14
```

```
Cont.:
```

```
Current A   Capacity W
----------- ------------
2.94        400.00
2.50        400.00
2.69        400.00
3.25        400.00
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
sensor_status	Sensor Status	3
power	Power W	4
voltage	Voltage V	5
current	Current A	6
capacity	Capacity W	7
temperature	Temperature	N/A
fan_speed	Fan Speed	N/A
fan_sensor_status	Fan Status	N/A
serial_number	Serial No	N/A
original_serial_number	Original Serial No	N/A
part_number	Part No	N/A
original_part_number	Original Part No	N/A
hw_rev	HW Rev	N/A
requires_service	Requires Service	N/A
service_reason	Service Reason	N/A
currently_functioning	Currently Functioning	N/A
switch_id	Switch ID	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing the configured InfiniBand switch BBUs

Use the **switch_bbu_list** command to list the configured InfiniBand switch BBUs.

```
switch_bbu_list [ switch_bbu=ComponentId ] [ switch=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
switch_bbu	IB switch BBU to list.	N	All IB switch BBUs
switch	IB switch whose BBUs are to be listed.	N	All IB switches

Example:

```
switch_bbu_list
```

Output:

```
Component ID          Status   Sensor Status   Remaining capacity   Full charged capacity
--------------------- -------- --------------- -------------------- ----------------------
```

```
1:IB_Switch_BBU:1:1   OK       OK              71520 mWh            80660 mWh
1:IB_Switch_BBU:1:2   OK       OK              46060 mWh            55520 mWh
1:IB_Switch_BBU:2:1   OK       OK              80800 mWh            80800 mWh
1:IB_Switch_BBU:2:2   OK       OK              77840 mWh            80120 mWh

Cont.:

Percent Charged   Charger Status   Calibration Status   Last Recondition Date
---------------- ---------------- -------------------- ----------------------
100%              Fully charged    N/A                  N/A
100%              Fully charged    N/A                  N/A
100%              Fully charged    N/A                  N/A
100%              Fully charged    N/A                  N/A
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
sensor_status	Sensor Status	3
relative_capacity	Remaining capacity	4
absolute_capacity	Full charged capacity	5
relative_capacity_percent	Percent Charged	6
charging_state	Charger Status	7
remaining_charging_time	Remaining Charging Time	N/A
calibration_status	Calibration Status	8
last_calibration_date	Last Recondition Date	9
fw_version	FW	N/A
serial_number	Serial No	N/A
original_serial_number	Original Serial No	N/A
part_number	Part No	N/A
original_part_number	Original Part No	N/A
voltage	Voltage	N/A
exp_voltage	Expected Voltage	N/A
temperature	Temperature	N/A
manufacture_date	Manufacture Date	N/A
designed_capacity	Designed Capacity	N/A
absolute_charge	Absolute Charge	N/A
test_status	Test Status	N/A
last_test_date	Last Test Date	N/A
can_charge	Charge	N/A
can_discharge	Discharge	N/A
charge_rate	Charge Rate	N/A
max_error	Max Error	N/A
hw_rev	HW Rev	N/A
requires_service	Requires Service	N/A
service_reason	Service Reason	N/A
currently_functioning	Currently Functioning	N/A
switch_id	Switch ID	N/A
active_alarms	Active Alarms	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing the configured InfiniBand switch fans

Use the **switch_fan_list** command to list the configured InfiniBand switch fans.

```
switch_fan_list [ switch_fan=ComponentId ] [ switch=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
switch_fan	IB switch fan to list.	N	All IB switch fans
switch	IB switch to list the BBUs of.	N	All IB switches

Example:

```
switch_fan_list
```

Output:

```
Component ID          Status   Speed      Sensor Status   Peer Speed   Peer Sensor Status
--------------------- -------- ---------- --------------- ------------ --------------------
1:IB_Switch_Fan:1:1   OK       10608.00   OK              8998.00      OK
1:IB_Switch_Fan:1:2   OK       10526.00   OK              8939.00      OK
1:IB_Switch_Fan:1:3   OK       10608.00   OK              8998.00      OK
1:IB_Switch_Fan:1:4   OK       10691.00   OK              9242.00      OK
1:IB_Switch_Fan:2:1   OK       10526.00   OK              9118.00      OK
1:IB_Switch_Fan:2:2   OK       10445.00   OK              9118.00      OK
1:IB_Switch_Fan:2:3   OK       10445.00   OK              8998.00      OK
1:IB_Switch_Fan:2:4   OK       10526.00   OK              8998.00      OK
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
speed	Speed	3
sensor_status	Sensor Status	4
peer_speed	Peer Speed	5
peer_sensor_status	Peer Sensor Status	6
serial_number	Serial No	N/A
original_serial_number	Original Serial No	N/A
part_number	Part No	N/A
original_part_number	Original Part No	N/A
hw_rev	HW Rev	N/A
requires_service	Requires Service	N/A
service_reason	Service Reason	N/A
currently_functioning	Currently Functioning	N/A
switch_id	Switch ID	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Access control commands

This section describes the command-line interface (CLI) for user access control.

- **Adding an access control definition**
  Use the **access_define** command to define an association between a user group and a host.
- **Deleting an access control definition**
  Use the **access_delete** command to delete an access control definition.
- **Listing access control definitions**
  Use the **access_list** command to list access control definitions.
- **Adding an LDAP server definition**
  Use the **ldap_add_server** command to add an LDAP server definition.
- **Testing an LDAP configuration**
  Use the **ldap_test** command to authenticate the specified user against an LDAP server, based on the existing configuration.
- **Listing LDAP configuration parameters**
  Use the **ldap_config_get** command to display system parameters that control user authentication against a specified LDAP server.
- **Configuring LDAP in the system**
  Use the **ldap_config_set** command to configure general system parameters that control user authentication against LDAP servers.
- **Listing LDAP servers defined in the system**
  Use the **ldap_list_servers** command to list LDAP servers defined in the system.
- **Listing LDAP server users**
  Use the **ldap_user_list** command to list LDAP server users.

- **Listing LDAP-based authentication mode**
  Use the **ldap_mode_get** command to list LDAP-based authentication mode.
- **Enabling or disabling LDAP-based authentication mode**
  Use the **ldap_mode_set** command to enable or disable LDAP-based authentication mode.
- **Updating an LDAP server definition**
  Use the **ldap_update_server** command to update an existing server configuration.
- **Removing an LDAP server definition**
  Use the **ldap_remove_server** command to remove an LDAP server definition.
- **Launching the ldapsearch utility**
  Use the **ldap_search** command to launch the ldapsearch utility.
- **Defining a new user**
  Use the **user_define** command to define a new user.
- **Deleting a user**
  Use the **user_delete** command to delete a user.
- **Adding users to user groups**
  Use the **user_group_add_user** command to add a user to a user group.
- **Creating user groups**
  Use the **user_group_create** command to create a user group.
- **Deleting a user group**
  Use the **user_group_delete** command to delete a user group.
- **Listing user groups**
  Use the **user_group_list** command to list all user groups or a specific one.
- **Removing a user from a user group**
  Use the **user_group_remove_user** command to remove a user from a user group.
- **Renaming user groups**
  Use the **user_group_rename** command to rename a user group.
- **Updating a user group**
  Use the **user_group_update** command to update a user group.
- **Listing users**
  Use the **user_list** command to list all users or a specific user.
- **Renaming users**
  Use the **user_rename** command to rename a user.
- **Updating a user definition**
  Use the **user_update** command to update a user definition.
- **Creating a new domain**
  Use the **domain_create** command to create a domain.
- **Updating a domain definition**
  Use the **domain_update** command to update a domain definition.
- **Renaming a domain**
  Use the **domain_rename** command to rename a domain.
- **Deleting a domain**
  Use the **domain_delete** command to delete a domain.
- **Listing domains**
  List all domains or the specified one.
- **Listing users per domain**
  Use the **domain_list_users** command to list users associated with domain(s).
- **Listing objects in domains**
  Use the **domain_list_objects** command to list objects attached to domain(s).
- **Listing the global domain**
  Use the **domain_global_list** to list the global domain.
- **Attaching an object to a domain**
  Use the **domain_attach_object** command to associate an object with a domain.
- **Disassociating object from a domain**
  Use the **domain_detach_object** command to disassociate object from a domain.
- **Associating users to a domain**
  Use the **domain_add_user** command to associate a user to a domain.
- **Removing a user from a domain**
  Use the **domain_remove_user** command to remove a user from a domain.
- **Adding a pool to a domain**
  Use the **domain_add_pool** command to add a pool to a domain.
- **Removing a pool from a domain**
  Use the **domain_remove_pool** command to remove a pool from a domain.
- **Moving a pool from one domain to another**
  Use the **domain_move_pool** command to move a pool from one domain to another.
- **Setting the domain attribute**
  Use the **domain_manage** command to set or clear the Managed attribute of a domain.
- **Setting domain related policies**
  Use the **domain_policy_set** command to set domain-related policies.
- **Displaying domain-related policies**
  Use the **domain_policy_get** command to display domain-related policies.
- **Specifying a user associated with IBM Hyper-Scale Manager**
  Use the **designate_msm_user_set** command to specify the name of the user that is associated with the IBM Hyper-Scale Manager.
- **Retrieving the user associated with the IBM Hyper-Scale Manager**
  Use the **designate_msm_user_get** command to retrieve the name of the user associated with the IBM Hyper-Scale Manager.
- **Setting the application administrator's scope of commands**
  Use the **appadmin_capabilities_set** command to define whether an Application Administrator is authorized to perform the basic or advanced set of

commands.

- **Getting the application administrator's scope of commands**
  Use the **appadmin_capabilities_get** command to display the state of the Application Administrator's capabilities.

Adding an access control definition

Use the **access_define** command to define an association between a user group and a host.

```
access_define user_group=UserGroup < host=HostName | cluster=ClusterName >
```

Parameters

Name	Type	Description	Mandatory
user_group	Object name	User group to be associated with the host or cluster.	Y
host	Object name	Host to be associated with the user group.	N
cluster	Object name	Cluster to be associated with the user group.	N

This command associates a user group with a host or a cluster. Hosts and clusters can only be associated with a single user group.

Example:

```
access_define host=host1 user_group=usergroup1
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- USER_GROUP_NAME_DOES_NOT_EXIST
  The user group name does not exist.
- CLUSTER_BAD_NAME
  The cluster name does not exist.
- HOST_BAD_NAME
  The host name does not exist.
- HOST_BELONGS_TO_CLUSTER
  This host already belongs to a cluster.

Deleting an access control definition

Use the **access_delete** command to delete an access control definition.

```
access_delete user_group=UserGroup < host=HostName | cluster=ClusterName >
```

Parameters

Name	Type	Description	Mandatory
user_group	Object name	The user group specified in the access control definition that should be deleted.	Y
host	Object name	The host specified in the access control definition that should be deleted.	N
cluster	Object name	The cluster specified in the access control definition that should be deleted.	N

This command deletes an association between the user group and host or cluster. The operation fails if no such access definition exists. When a host is removed from a cluster, the host's associations become the cluster's associations. This allows a continued mapping of operations, so that all scripts continue to work.

Example:

```
access_delete user_group=usergroup1
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- USER_GROUP_NAME_DOES_NOT_EXIST
  The user group name does not exist.
- USER_GROUP_DOES_NOT_HAVE_ACCESS_TO_CLUSTER
  User Group does not have access to cluster.
- CLUSTER_BAD_NAME
  The cluster name does not exist.
- HOST_BAD_NAME
  The host name does not exist.
- HOST_BELONGS_TO_CLUSTER
  This host already belongs to a cluster.
- USER_GROUP_DOES_NOT_HAVE_ACCESS_TO_HOST
  User Group does not have access to host.

Listing access control definitions

Use the **access_list** command to list access control definitions.

```
access_list [ user_group=UserGroup ] [ host=HostName | cluster=ClusterName ]
```

Parameters

Name	Type	Description	Mandatory	Default
user_group	Object name	Filters the access control listing to display only this user group.	N	All user groups.
host	Object name	Filters the access control listing to display only this host.	N	All hosts.
cluster	Object name	Filters the access control listing to display only this cluster.	N	All clusters.

The list can be displayed for all access control definitions, or it can be filtered for a specific user group, host/cluster, or both.

Field ID	Field output	Default position
type	Type	1
name	Name	2
user_group	User Group	3

Example:

```
access_list host=buffyvam
```

Output:

```
Type   Name       User Group
host   buffyvam   testing
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Return codes

- HOST_BAD_NAME
  The host name does not exist.
- CLUSTER_BAD_NAME
  The cluster name does not exist.
- USER_GROUP_NAME_DOES_NOT_EXIST
  The user group name does not exist.

Adding an LDAP server definition

Use the **ldap_add_server** command to add an LDAP server definition.

```
ldap_add_server  fqdn=Fqdn  [ address=Address ] base_dn=LdapDn [ certificate=PemCertificate ] [ port=PortNum ]
 [ secure_port=PortNum ]
```

Parameters

Name	Type	Description	Mandatory	Default
fqdn	N/A	FQDN of the LDAP server.	Y	N/A
address	N/A	IP address of the LDAP server. Only required when DNS is not available for FQDN to IP address resolution.	N	none
base_dn	N/A	Base_DN of the LDAP server. Serves as the starting reference point for searches.	Y	N/A
certificate	N/A	The public certificate or certificate chain of the LDAP server to be added (see below for details).	N	no certificate
port	Integer	The port number.	N	389
secure_port	Integer	The secure port number.	N	636

The certificate parameter

The value of the certificate parameter is the content of a PEM file with asterisks instead of newlines. Chained certificates are supported. The total maximal length of a PEM file holding chained certificates (leaf first, root last) is 15360 characters (including the asterisk characters). In Windows, you can drag-and-drop a PEM file from the Windows Explorer to the appropriate location in the XCLI session window; the content will be added automatically.

Example:

```
ldap_add_server fqdn=ldap.example.com address=192.0.2.1
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- MAX_LDAP_SERVERS_REACHED
  The number of defined LDAP servers already reached the maximum.
- ADDRESS_CURRENTLY_ASSOCIATED_WITH_ANOTHER_LDAP_SERVER
  The specified IP address is currently associated with another LDAP server.
- LDAP_SERVER_EXISTS
  LDAP server with the specified FQDN already exists.
- SSL_CERTIFICATE_CHAIN_EMPTY
  No certificates were found in the input.
- SSL_CERTIFICATE_HAS_EXPIRED
  The SSL certificate has expired.
- SSL_CERTIFICATE_INVALID_FORMAT
  The SSL certificate format is invalid or corrupted.

- SSL_CERTIFICATE_ISSUER_NOT_FOUND
  The SSL certificate issuer was not found in the certificate chain.
- SSL_CERTIFICATE_NOT_YET_VALID
  The SSL certificate is not yet valid.
- SSL_CERTIFICATE_VERIFICATION_FAILED
  The SSL certificate chain verification failed.
- SSL_CERTIFICATE_VERIFICATION_INTERNAL_ERROR
  The SSL certificate verification has failed because of an internal system error.

Testing an LDAP configuration

Use the **ldap_test** command to authenticate the specified user against an LDAP server, based on the existing configuration.

```
ldap_test [ fqdn=Fqdn ] user=UserName    password=Password
```

Parameters

Name	Type	Description	Mandatory	Default
fqdn	N/A	FQDN of an LDAP server.	N	All servers
user	String	The username of the tested user.	Y	N/A
password	String	The password of the tested user.	Y	N/A

Example:

```
ldap_test user=user1 password=pass1
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- NO_LDAP_SERVERS_ARE_DEFINED
  No LDAP servers are defined in the system
- LDAP_SERVER_NOT_DEFINED
  LDAP server *Server FQDN* is not defined in the system.
- LDAP_IS_NOT_FULLY_CONFIGURED
  LDAP is not fully configured.

  Troubleshooting: Check LDAP settings.
- NO_LDAP_SERVERS_WITH_CERTIFICATE_ARE_DEFINED
  No LDAP servers with an LDAP certificate are defined in the system.
- SSL_CERTIFICATE_HAS_EXPIRED_FOR_SERVER
  SSL certificate of LDAP server '*Server FQDN*' expired on *Expiration Date*.
- USER_IS_PREDEFINED_IN_THE_SYSTEM
  The user is predefined in the system.
- LOGIN_FAILURE_USER_CANNOT_BE_UNIQUELY_AUTHENTICATED_BY_LDAP_SERVER
  User *User Name* was not uniquely authenticated by LDAP server '*Server FQDN*'.
- LOGIN_FAILURE_LDAP_SERVER_UNREACHABLE
  No LDAP server can be reached.
- LOGIN_FAILURE_XIV_USER_NOT_AUTHENTICATED_BY_LDAP_SERVER
  XIV User '*XIV User*' was not authenticated by LDAP server '*Server FQDN*'.
- LOGIN_FAILURE_LDAP_SERVER_UNREACHABLE_OR_USER_NOT_FOUND
  User *User Name* was not found in LDAP servers '*Servers FQDN*'.

- LOGIN_FAILURE_INVALID_BASE_DN
  The base dn of server '*Server FQDN*' is invalid.
- LOGIN_FAILURE_USER_NOT_AUTHENTICATED_BY_LDAP_SERVER
  User *User Name* was not authenticated by LDAP server '*Server FQDN*'.
- LOGIN_FAILURE_USER_HAS_NO_RECOGNIZED_ROLE
  User *User Name* has no recognized LDAP role.
- LOGIN_FAILURE_USER_HAS_MORE_THAN_ONE_RECOGNIZED_ROLE
  User *User Name* has more than one recognized LDAP role.
- LOGIN_FAILURE_USER_MISSING_ID_ATTRIBUTE
  User *User Name* is missing the LDAP ID attribute '*Attribute*'.
- LOGIN_FAILURE_USER_MISSING_GROUP_ATTRIBUTE
  User *User Name* is missing the group attribute '*Attribute*'.
- LOGIN_FAILURE_USER_NOT_FOUND_IN_LDAP_SERVERS
  User *User Name* was not found in LDAP servers.
- LDAP_ROLE_UNRECOGNIZED
  The LDAP role for the user is not recognized in the system.
- LDAP_SERVER_NOT_FOUND
  No LDAP server with the specified FQDN is defined in the system.
- LDAP_AUTHENTICATION_IS_NOT_ACTIVE
  LDAP authentication is not active.

Listing LDAP configuration parameters

Use the **ldap_config_get** command to display system parameters that control user authentication against a specified LDAP server.

```
ldap_config_get
```

A successful execution of this command depends on connecting to a valid LDAP server.

The output of the command does not list LDAP servers. For the list of LDAP servers, use the ldap_list_servers command.

The xiv_password parameter is not listed.

Example:

```
ldap_config_get
```

Output:

```
Name                       Value
current_server
base_dn
version                    3
xiv_group_attrib
storage_admin_role
read_only_role
session_cache_period       20
bind_time_limit            20
user_id_attrib             objectSiD
first_expiration_event     30
second_expiration_event    14
third_expiration_event     7
use_ssl                    no
xiv_user
```

Field ID	Field output	Default position
name	Name	1
value	Value	2

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Configuring LDAP in the system

Use the **ldap_config_set** command to configure general system parameters that control user authentication against LDAP servers.

```
ldap_config_set [ user_name_attrib=LdapAttrib ] [
        xiv_group_attrib=LdapAttrib ]
 [ storage_admin_role=LdapRole ] [
        read_only_role=LdapRole ]
 [ security_admin_role=LdapRole ] [
        storage_integration_admin_role=LdapRole ]
 [
        xiv_host_profiler_role=LdapRole ] [ use_ssl=<yes|no> ]
 [
        user_id_attrib=LdapAttrib ] [ session_cache_period=Minutes ]
 [
        bind_time_limit=Seconds ] [ first_expiration_event=Days ]
 [
        second_expiration_event=Days ] [ third_expiration_event=Days ]
 [
        version=LdapVersion ] [ xiv_user=LdapAttrib ] [
        xiv_password=LdapAttrib ]
 [ server_type=<SUN DIRECTORY|MICROSOFT
        ACTIVE DIRECTORY|OPEN LDAP> ]
 [ group_search_depth=Depth ] [
        group_search_max_queries=Number ]
 [
        group_search_stop_when_found=<yes|no> ]
```

Parameters

Name	Type	Description	Mandatory	Default
user_name_attrib	String	User name attribute for queries. If not specified, it is set to **uid** for SUN Directory servers and **userPrincipalName** for Microsoft Active Directory servers.	N	According to server type
xiv_group_attrib	String	LDAP attribute designated to hold system-mapped roles.	N	none
storage_admin_role	String	LDAP value mapped to the Storage Administrator role. Multiple (up to 8) values are supported and must be separated with a semicolon (;). Multiple roles are not available for SUN Directory LDAP Servers. Up to 499 characters are allowed.	N	none
read_only_role	String	LDAP value mapped to the Read Only role. Multiple (up to 8) values are supported and must be separated with a semicolon (;). Up to 499 characters are allowed.	N	none
security_admin_role	String	LDAP value mapped to the Security Administrator role. Multiple (up to 8) values are supported and must be separated with a semicolon (;). Up to 499 characters are allowed.	N	none
storage_integration_admin_role	String	LDAP value mapped to the Storage Integration Administrator role. Multiple (up to 8) values are supported and must be separated with a semicolon (;). Up to 499 characters are allowed.	N	none
xiv_host_profiler_role	String	LDAP value mapped to the xiv host profiler role. Multiple (up to 8) values are supported and must be separated with a semicolon (;). Up to 499 characters are allowed.	N	none
use_ssl	Boolean	Indicates whether secure LDAP is mandatory.	N	no
user_id_attrib	String	The LDAP attribute set to identify the user (in addition to user DN) when recording user operations in the event log.	N	objectSiD
session_cache_period	Positive integer	Duration of keeping user credentials before attempting to re-login the user.	N	20
bind_time_limit	Positive integer	The duration after which the next LDAP server on the LDAP server list will be called.	N	0. If set to the default, the LDAP server is contacted for every command. Performance issues depend on its availability.
first_expiration_event	Positive integer	The number of days before the expiration of certificate, when the first alert is issued (severity: warning).	N	30/14/7 (third is smallest)
second_expiration_event	Positive integer	The number of days before the expiration of certificate, when the second alert is issued (severity: warning).	N	30/14/7 (third is smallest)
third_expiration_event	Positive integer	The number of days before the expiration of certificate, when the third alert is issued (severity: warning).	N	30/14/7 (third is smallest)
version	Positive integer	Version of LDAP used (only version 3 is supported).	N	3
xiv_user	String	The user for LDAP queries.	N	none
xiv_password	String	The password of user for LDAP queries.	N	none
server_type	Enumeration	Type of the LDAP server.	N	none
group_search_depth	Positive integer	The depth of group hierarchy to search in.	N	0

Name	Type	Description	Mandatory	Default
group_search_max_queries	Positive integer	Maximum number of group queries to perform per server.	N	39
group_search_stop_when_found	Boolean	Stop the group search when a group match is found.	N	yes

LDAP access permissions are not enforced for predefined users. These predefined users are authenticated by the IBM storage system and not by LDAP even if LDAP authentication is enabled.

Predefined user names are:

- admin
- technician
- xiv_development
- xiv_maintenance

When an LDAP user, whose user name is identical with a predefined name, attempts to log into the system with LDAP authentication enabled, access will normally be denied, because:

- the user is not authenticated against LDAP, but rather against the storage system
- the user's (LDAP) password most likely does not match the storage system password.

However, if the user attempts to log into the system using the password of the corresponding predefined user, he or she will be granted the rights of the corresponding predefined user regardless of LDAP settings (for example, the user's association with the Application Administrator role), because LDAP authentication for predefined users is not required.

When an LDAP user, with a user name identical to any of the predefined names, tries to log into the system (when LDAP authentication is enabled), that user will normally be denied access, since he or she is not authenticated against LDAP (but rather against the storage system), and the user's (LDAP) password most likely does not match the storage system password. However, logging into the system using the password of the corresponding predefined user, the user will be granted the rights of the corresponding predefined user, regardless of LDAP settings (for example, association with an Application Administrator role), because LDAP is not consulted for predefined users.

Example:

```
ldap_config_set
            storage_admin_role="CN=EXMPLEAdmins,OU=EXMPLELab,DC=CROSSQA,DC=COM"
            re
ad_only_role="CN=EXMPLEOpers,OU=EXMPLELab,DC=CROSSQA,DC=COM"
            user_id_attrib=sAMA
ccountName user_name_attrib=sAMAccountName
            xiv_group_attrib=memberOf

 xiv_user="CN=allmighty,CN=Users,DC=CROSSQA,DC=COM"
            xiv_password=the_password

          security_admin_role="CN=EXMPLESecAdmins,OU=EXMPLELab,DC=CROSSQA,DC=COM"
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- LDAP_IS_NOT_FULLY_CONFIGURED
  LDAP is not fully configured.

  Troubleshooting: Check LDAP settings.

- LDAP_CONFIG_CHANGE_IS_ILLEGAL_WHEN_AUTHENTICATION_IS_ACTIVE
  This LDAP configuration change is invalid when LDAP-based authentication is active.

  Troubleshooting: Disable the LDAP-based authentication and then change the LDAP configuration.

- LDAP_ROLE_ALREADY_USED
  This LDAP role is already in use in the LDAP configuration or in a user group.

- NO_LDAP_SERVERS_WITH_CERTIFICATE_ARE_DEFINED
  No LDAP servers with an LDAP certificate are defined in the system.

- INVALID_EXPIRATION_EVENT_DATES
  The dates for expiration events must be in ascending order.

- LDAP_READ_ONLY_ROLE_HAS_WRONG_NUMBER_OF_PARTS
  Too many parts in the LDAP read-only role.

  Troubleshooting: Role parts are divided by the ';' symbol. The number of parts should be between 0 and 8.

- LDAP_ROLE_HAS_DUPLICATED_PARTS
  The LDAP role contains duplicated parts.

- LDAP_STORAGE_ADMIN_ROLE_HAS_WRONG_NUMBER_OF_PARTS
  Too many parts in the LDAP storage administrator role.

  Troubleshooting: Role parts are divided by the ';' symbol. The number of parts should be between 0 and 8.

- LDAP_SECURITY_ADMIN_ROLE_HAS_WRONG_NUMBER_OF_PARTS
  Too many parts in the LDAP security administrator role.

  Troubleshooting: Role parts are divided by the ';' symbol. The number of parts should be between 0 and 8.

- LDAP_STORAGE_INTEGRATION_ADMIN_ROLE_HAS_WRONG_NUMBER_OF_PARTS
  Too many parts in the LDAP storage integration administrator role.

  Troubleshooting: Role parts are divided by the ';' symbol. The number of parts should be between 0 and 8.

- LDAP_XIV_HOST_PROFILER_ROLE_HAS_WRONG_NUMBER_OF_PARTS
  Too many parts in the LDAP XIV host profiler role.

  Troubleshooting: Role parts are divided by the ';' symbol. The number of parts should be between 0 and 8.

Listing LDAP servers defined in the system

Use the **ldap_list_servers** command to list LDAP servers defined in the system.

```
ldap_list_servers  [ fqdn=Fqdn ]
```

Parameters

Name	Description	Mandatory	Default
fqdn	FQDN of a specific server to list.	N	All servers.

This command lists the LDAP servers defined in the system along with their type description and the indication whether they are mandatory.

Example:

```
ldap_list_servers fqdn
```

Output:

```
<code value="SUCCESS"/>
        <empty_table_message value="No LDAP servers are defined in the system"/>
        <last_change_index value="367896"/>
        <status value="0"/>
        <status_str value="Command completed successfully"/>
```

Field ID	Field output	Default position
fqdn	FQDN	1
address	Address	2
base_dn	Base DN	3
has_certificate	Has Certificate	4
expiration_date	Expiration Date	5
valid_certificate	Valid Certificate	N/A
accessible	Accessible	N/A
port	Port	6
secure_port	Secure Port	7

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Listing LDAP server users

Use the **ldap_user_list** command to list LDAP server users.

```
ldap_user_list role=Category [ domain=DomainName ]
```

Parameters

Name	Type	Description	Mandatory	Default
role	Enumeration	The role of the users to be retrieved from the LDAP server. The available roles are: storageadmin and readonly.	Y	N/A
domain	Object name	The domain name.	N	All Domains

This command retrieves a list of users from the LDAP server by a specific role.

Field ID	Field output	Default position
user_name	User Name	1
user_role	Role	2

Example:

```
ldap_user_list role=storageadmin
```

Output:

```
User Name        Role
---------------- -----------
readonly_user    Read Only
restldapread     Read Only
test_readonly    Read Only
xivreadonly      Read Only
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- LDAP_AUTHENTICATION_IS_NOT_ACTIVE
  LDAP authentication is not active.
- LDAP_IS_NOT_FULLY_CONFIGURED
  LDAP is not fully configured.

  Troubleshooting: Check LDAP settings.
- NO_LDAP_SERVERS_ARE_DEFINED
  No LDAP servers are defined in the system
- NO_LDAP_SERVERS_WITH_CERTIFICATE_ARE_DEFINED
  No LDAP servers with an LDAP certificate are defined in the system.
- LOGIN_FAILURE_XIV_USER_NOT_AUTHENTICATED_BY_LDAP_SERVER
  XIV User '*XIV User*' was not authenticated by LDAP server '*Server FQDN*'.
- LOGIN_FAILURE_LDAP_SERVER_UNREACHABLE
  No LDAP server can be reached.
- LOGIN_FAILURE_INVALID_BASE_DN
  The base dn of server '*Server FQDN*' is invalid.

Listing LDAP-based authentication mode

Use the **ldap_mode_get** command to list LDAP-based authentication mode.

```
ldap_mode_get
```

The command succeeds regardless of whether the LDAP server is accessible.

Example:

```
ldap_mode_get
```

Output:

```
Mode
Inactive
```

Field ID	Field output	Default position
mode	Mode	1

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Enabling or disabling LDAP-based authentication mode

Use the **ldap_mode_set** command to enable or disable LDAP-based authentication mode.

```
ldap_mode_set  mode=Mode
```

Parameters

Name	Type	Description	Mandatory
mode	Boolean	The required state of LDAP authentication. Available values: Active, Inactive.	Y

Example:

```
ldap_mode_set mode=active
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- USERS_ASSOCIATED_TO_DOMAIN_ARE_YOU_SURE_YOU_WANT_TO_ENABLE_LDAP_AUTHENTICATION
  There are users associated with domains. Are you sure you want to enable LDAP authentication?
- ARE_YOU_SURE_YOU_WANT_TO_ENABLE_LDAP_AUTHENTICATION
  Are you sure you want to enable LDAP authentication?
- ARE_YOU_SURE_YOU_WANT_TO_DISABLE_LDAP_AUTHENTICATION
  Are you sure you want to disable LDAP authentication?

Return codes

- LDAP_IS_NOT_FULLY_CONFIGURED
  LDAP is not fully configured.

  Troubleshooting: Check LDAP settings.
- NO_LDAP_SERVERS_WITH_CERTIFICATE_ARE_DEFINED
  No LDAP servers with an LDAP certificate are defined in the system.

- NO_LDAP_SERVERS_ARE_DEFINED
  No LDAP servers are defined in the system

Updating an LDAP server definition

Use the **ldap_update_server** command to update an existing server configuration.

```
ldap_update_server  fqdn=Fqdn [ address=Address ] [ base_dn=LdapDn ]  [ port=PortNum ] [ secure_port=PortNum ]
   [ certificate=PemCertificate  | remove_certificate=<no|yes> ]
```

Parameters

Name	Type	Description	Mandatory	Default
fqdn	N/A	FQDN of the LDAP server.	Y	N/A
address	N/A	IP address of the LDAP server.	N	none
certificate	N/A	The public certificate or certificate chain of the LDAP server to be added (see below for details).	N	no certificate
remove_certificate	Boolean	Defines whether to remove the certificate.	N	no
base_dn	N/A	Base_DN of the LDAP directory.	N	none
port	Integer	The port number.	N	none
secure_port	Integer	The secure port number.	N	none

The certificate parameter:

The value of the certificate parameter is the content of a PEM file with asterisks instead of newlines. Chained certificates are supported. The total maximal length of a PEM file holding chained certificates (leaf first, root last) is 15360 characters (including the asterisk characters). In Windows, you can drag-and-drop a PEM file from the Windows Explorer to the appropriate location in the XCLI session window; the content will be added automatically.

Example:

```
ldap_update_server fqdn=ldap.example.com address=192.0.2.1
        remove_certificate=yes
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_UPDATE_LDAP_SERVER
  Are you sure you want to update the LDAP server configuration?

Return codes

- LDAP_SERVER_NOT_FOUND
  No LDAP server with the specified FQDN is defined in the system.
- ADDRESS_CURRENTLY_ASSOCIATED_WITH_ANOTHER_LDAP_SERVER
  The specified IP address is currently associated with another LDAP server.
- NO_UPDATE_PARAMETERS_SPECIFIED
  No LDAP server parameters were specified for the update.
- SSL_CERTIFICATE_CHAIN_EMPTY
  No certificates were found in the input.
- SSL_CERTIFICATE_HAS_EXPIRED
  The SSL certificate has expired.
- SSL_CERTIFICATE_INVALID_FORMAT
  The SSL certificate format is invalid or corrupted.

- SSL_CERTIFICATE_ISSUER_NOT_FOUND
  The SSL certificate issuer was not found in the certificate chain.
- SSL_CERTIFICATE_NOT_YET_VALID
  The SSL certificate is not yet valid.
- SSL_CERTIFICATE_VERIFICATION_FAILED
  The SSL certificate chain verification failed.
- SSL_CERTIFICATE_VERIFICATION_INTERNAL_ERROR
  The SSL certificate verification has failed because of an internal system error.

Removing an LDAP server definition

Use the **ldap_remove_server** command to remove an LDAP server definition.

```
ldap_remove_server  fqdn=Fqdn
```

Parameters

Name	Description	Mandatory
fqdn	FQDN of the server to remove.	Y

Example:

```
ldap_remove_server fqdn=cloud.example.com
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_REMOVE_LDAP_SERVER
  Are you sure you want to remove the LDAP server?

Return codes

- LDAP_SERVER_NOT_FOUND
  No LDAP server with the specified FQDN is defined in the system.
- LDAP_IS_ACTIVE_BUT_THIS_IS_THE_LAST_SERVER
  Deleting the last LDAP server is illegal when LDAP authentication is active.
- LDAP_USES_SSL_BUT_THIS_IS_THE_LAST_SERVER_WITH_CERTIFICATE
  Deleting the last LDAP server with a valid SSL certificate is illegal when LDAP authentication is active and uses SSL.

Launching the ldapsearch utility

Use the **ldap_search** command to launch the ldapsearch utility.

```
ldap_search  user=UserName fqdn=Fqdn [ second_cmd=<yes|no> password=Password ]
```

Parameters

Name	Type	Description	Mandatory	Default
user	Object name	The username to search for.	Y	N/A
password	N/A	The user password to search for.	N	empty
second_cmd	Boolean	Defines whether to invoke the second lsearch command.	N	no
fqdn	N/A	FQDN of LDAP server to query.	Y	N/A

There are 2 LDAP search commands executed in the authentication process. The second one can be issued by setting the second_cmd parameter to yes.

Example:

```
ldap_search fqdn user password
```

Output:

```
Name            Index   Value
--------------- ------- ------------------------------------------------------
command_line    0       ldapsearch -H ldap://ldapwin2003.example.com:389...
returncode      0       0
stderr          0
stdout          0       dn: CN=employee,CN=Users,DC=xivldap2,DC=com
stdout          1       description: Group One
stdout          2       objectSid:: AQUAAAAAAAUVAAAAYcKhSnhmt01IPSuAbQQAAA==
stdout          3
stdout          4
```

Field ID	Field output	Default position
name	Name	1
index	Index	2
value	Value	3

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- LOGIN_FAILURE_USER_NOT_AUTHENTICATED_BY_LDAP_SERVER
  User *User Name* was not authenticated by LDAP server '*Server FQDN*'.
- LOGIN_FAILURE_USER_MISSING_GROUP_ATTRIBUTE
  User *User Name* is missing the group attribute '*Attribute*'.
- LDAP_SERVER_NOT_FOUND
  No LDAP server with the specified FQDN is defined in the system.
- LOGIN_FAILURE_LDAP_SERVER_UNREACHABLE
  No LDAP server can be reached.
- LDAP_SERVER_NOT_DEFINED
  LDAP server *Server FQDN* is not defined in the system.
- LDAP_ROLE_UNRECOGNIZED
  The LDAP role for the user is not recognized in the system.
- LOGIN_FAILURE_USER_HAS_NO_RECOGNIZED_ROLE
  User *User Name* has no recognized LDAP role.
- LOGIN_FAILURE_USER_CANNOT_BE_UNIQUELY_AUTHENTICATED_BY_LDAP_SERVER
  User *User Name* was not uniquely authenticated by LDAP server '*Server FQDN*'.
- LOGIN_FAILURE_XIV_USER_NOT_AUTHENTICATED_BY_LDAP_SERVER
  XIV User '*XIV User*' was not authenticated by LDAP server '*Server FQDN*'.
- LOGIN_FAILURE_USER_HAS_MORE_THAN_ONE_RECOGNIZED_ROLE
  User *User Name* has more than one recognized LDAP role.
- LOGIN_FAILURE_USER_MISSING_ID_ATTRIBUTE
  User *User Name* is missing the LDAP ID attribute '*Attribute*'.
- USER_IS_PREDEFINED_IN_THE_SYSTEM
  The user is predefined in the system.
- LOGIN_FAILURE_INVALID_BASE_DN
  The base dn of server '*Server FQDN*' is invalid.
- LDAP_AUTHENTICATION_IS_NOT_ACTIVE
  LDAP authentication is not active.

Defining a new user

Use the **user_define** command to define a new user.

```
user_define user=UserName password=Password password_verify=Password category=Category [ email_address=email ]
 [ area_code=AreaCode number=PhoneNumber ] [ domain=DomainList [ exclusive=<yes|no> ] ]
```

Parameters

Name	Type	Description	Mandatory	Default
user	Object name	User name. User names are lower case.	Y	N/A
password	N/A	Password of the user to be created. The password must have between 6 and 16 characters. Any symbols are allowed, except the following: • double quotation (") • single quotation or apostrophe (') • grave accent (`) Passwords are case sensitive.	Y	N/A
password_verify	N/A	Password verification, which must be equal to the value of **password**.	Y	N/A
category	Enumeration	The role of the user to be created. Available roles: • storageadmin • applicationadmin • securityadmin • readonly • storageintegrationadmin	Y	N/A
email_address	N/A	Email address of this user. The email address specified here can be used for event notification. Entering this address is optional. The email address format is any legal email address.	N	N/A
number	N/A	Cellular phone number of the user for event notification via SMS, excluding the area code. Phone numbers and area codes can be a maximum of 63 digits, dashes (-) and periods (.)	N	N/A
area_code	N/A	Area code of the cellular phone number of the user. Phone numbers and area codes can be a maximum of 63 digits, dashes (-) and periods (.)	N	N/A
domain	N/A	The cluster will be attached to the specified domains. To specify more than one domain, separate them with a comma. To specify all existing domains, use "*".	N	none
exclusive	Boolean	Use yes to restrict the user to domain's objects.	N	yes

Email address and phone number are optional and can be used for event notification. The category (user role) may be only one of those specified above (other categories contain only a single predefined user).

The maximum number of users is 128.

There are two system-wide predefined users:

User name	User role
admin	Storage administrator
technician	Technician

Example:

```
user_define user=xiv_user1 password=s0mePassw0rd password_verify=s0mePassw0rd category=applicationadmin
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- USER_NAME_ALREADY_EXISTS
  The user name already exists.
- MAX_USERS_REACHED

The number of defined users already reached the maximum.

- PASSWORDS_DO_NOT_MATCH
  Make sure that passwords are identical.
- USER_PHONE_NUMBER_MUST_ACCOMPANY_AREA_CODE
  The phone number must be indicated together with the area code.
- LDAP_AUTHENTICATION_IS_ACTIVE
  This command is not available while LDAP authentication is active.
- DOMAIN_DOESNT_EXIST
  The domain does not exist.
- NON_EXCLUSIVE_USER_NOT_SECURITY_ADMIN
  Only a security administrator can add a user to a domain non-exclusively.
- USER_CANNOT_BE_ADDED_TO_A_DOMAIN
  The specified user cannot be associated with a domain.
- SIA_MUST_BE_ASSOCIATED_WITH_A_DOMAIN
  The storage integration administrator must be associated with a domain.

Deleting a user

Use the **user_delete** command to delete a user.

```
user_delete user=UserName
```

Parameters

Name	Type	Description	Mandatory
user	Object name	User to be deleted.	Y

Existing objects created by this user will retain an empty user reference after the user has been deleted.

Two predefined users are set system-wide: Admin and Technician. Predefined users cannot be deleted or renamed.

Example:

```
user_delete user=user1
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- LDAP_AUTHENTICATION_IS_ACTIVE
  This command is not available while LDAP authentication is active.
- USER_NAME_DOES_NOT_EXIST
  The user name does not exist.
- USER_CANNOT_BE_DELETED
  The user cannot be deleted.
- USER_IS_REFERRED_TO_BY_DEST
  The user is referred to by an event destination and therefore cannot be deleted.
- USER_OWNS_RECOVERY_KEY
  The user owns a recovery key and therefore cannot be deleted or renamed.
- REMOVAL_WOULD_CREATE_UNRESOLVABLE_REFERENCE_BETWEEN_USER_AND_USERGROUP
  Completing this operation will result in a user referring to a user group that is not in its domain.

  Troubleshooting: Remove the reference explicitly and re-run the command.

Adding users to user groups

Use the **user_group_add_user** command to add a user to a user group.

```
user_group_add_user user_group=UserGroup user=UserName
```

Parameters

Name	Type	Description	Mandatory
user_group	Object name	User group into which the user is to be added.	Y
user	Object name	User to be added to the user group.	Y

A user group can contain up to eight users.

A user may belong to only one user group.

Only users defined as Application Administrators can be assigned to a group.

This command fails when the user already belongs to the user group.

Example:

```
user_group_add_user user_group=ug1 user=user1
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- USER_GROUP_NAME_DOES_NOT_EXIST
  The user group name does not exist.
- USER_NAME_DOES_NOT_EXIST
  The user name does not exist.
- USER_ALREADY_INCLUDED_IN_ANOTHER_GROUP
  The user is included into another user group.
- USER_GROUP_ALREADY_INCLUDES_USER
  The user group already includes a user.
- ONLY_APPLICATION_ADMIN_USERS_CAN_BE_GROUPED
  User groups can only contain application administrators.
- USER_GROUP_HAS_MAXIMUM_NUMBER_OF_USERS
  The user group already contains the maximum number of users.
- LDAP_AUTHENTICATION_IS_ACTIVE
  This command is not available while LDAP authentication is active.

Creating user groups

Use the **user_group_create** command to create a user group.

```
user_group_create user_group=UserGroup [ access_all=<yes|no> ]   [ ldap_role=LdapRole ] [ domain=DomainList ]
```

Parameters

Name	Type	Description	Mandatory	Default
user_group	Object name	Name of the user group to be created.	Y	N/A

Name	Type	Description	Mandatory	Default
access_all	Boolean	Allows application administrators to perform their specified operations on all volumes and not just on a subset of the specific volumes.	N	no
ldap_role	String	The value representing the user group in LDAP.	N	[none]
domain	N/A	The user_group will be attached to the specified domains. To specify more than one domain, separate them with a comma. To specify all the existing domains, use "*".	N	none

A user group is a group of application administrators who share the same set of snapshot creation limitations. After user groups are created, the limitations of all the users in a user group can be updated with a single command. These limitations are enforced by associating the user groups with hosts or clusters.

Storage administrators create user groups and control the various application administrator's permissions. Hosts and clusters can be associated with only a single user group. When a user belongs to a user group that is associated with a host, it is possible to manage snapshots of the volumes mapped to that host.

User groups have the following limitations:

- Only users who are defined as application administrators can be assigned to a group.
- A user can belong to only a single user group.
- A user group can contain up to eight users.

User and host associations have the following properties:

- User groups can be associated with both hosts and clusters. This allows limiting application administrator access to specific volumes.
- A host that is part of a cluster cannot also be associated with a user group.
- When a host is added to a cluster the host's associations are broken. Limitations on the management of volumes mapped to the host is controlled by the cluster's association.
- When a host is removed from a cluster, the host's associations become the cluster 's associations, this allows continued mapping of operations so that all scripts continue to work.

Application administrator access level:

- The access_all parameter can be specified for application administrators only. When it is specified, it means that the user has an application administrator access level to all volumes, and can perform operations on all volumes and not just on a subset of the specific volume.

Example:

```
user_group_create user_group=ug1 ldap_role="App Admin 1" access_all=yes
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- USER_GROUP_NAME_ALREADY_EXISTS
  The user group name already exists.
- MAX_USER_GROUPS_REACHED
  The number of defined user groups already reached the maximum.
- LDAP_ROLE_ALREADY_USED
  This LDAP role is already in use in the LDAP configuration or in a user group.
- DOMAIN_DOESNT_EXIST
  The domain does not exist.

Deleting a user group

Use the **user_group_delete** command to delete a user group.

```
user_group_delete user_group=UserGroup
```

Parameters

Name	Type	Description	Mandatory
user_group	Object name	User group to be deleted.	Y

A user group can be deleted, even when it is associated with hosts or clusters. It can be deleted while in LDAP Authentication mode.

A user group can be deleted, even when it contains users. Deleting the user group does not delete the users contained in this group.

Example:

```
user_group_delete user_group=ug1
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_DELETE_LDAP_USER_GROUP
  One or more LDAP users might be associated with this user group. Are you sure you want to delete the user group?
- ARE_YOU_SURE_YOU_WANT_TO_DELETE_USER_GROUP
  Are you sure you want to delete the user group?
- ARE_YOU_SURE_YOU_WANT_TO_DELETE_POPULATED_USER_GROUP
  One or more internal users are associated with this user group. Are you sure you want to delete the user group?

Return codes

- USER_GROUP_NAME_DOES_NOT_EXIST
  The user group name does not exist.

Listing user groups

Use the **user_group_list** command to list all user groups or a specific one.

```
user_group_list [ user_group=UserGroup ] [ domain=DomainName ]
```

Parameters

Name	Type	Description	Mandatory	Default
user_group	Object name	The user group to be listed.	N	All user groups.
domain	Object name	The domain name.	N	All Domains

All the users included in the user group are listed.

Field ID	Field output	Default position
name	Name	1
access_all	Access All	2
ldap_role	LDAP Role	3
users	Users	4
creator	Creator	N/A

Example:

```
user_group_list
```

Output:

```
Name         Access All   LDAP Role   Users
myug1        yes          Group1
myOtherUG    yes          OtherGroup
ug1          yes          App Admin 1
ug2          yes          App Admin 2
```

Access control

User Category	Permission

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Removing a user from a user group

Use the **user_group_remove_user** command to remove a user from a user group.

```
user_group_remove_user user_group=UserGroup user=UserName
```

Parameters

Name	Type	Description	Mandatory
user_group	Object name	User group.	Y
user	Object name	User to be removed.	Y

This command fails when the user does not belong to the user group.

Deleting the user group's mapping is done by removing the role association. The user group itself is not deleted.

Example:

```
user_group_remove_user user_group=ug1 user=user1
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_REMOVE_USER
  Are you sure you want to remove the user from the user group?

Return codes

- USER_GROUP_NAME_DOES_NOT_EXIST
  The user group name does not exist.
- USER_NAME_DOES_NOT_EXIST
  The user name does not exist.
- USER_GROUP_DOES_NOT_INCLUDE_USER
  The user group does not include any user.
- LDAP_AUTHENTICATION_IS_ACTIVE
  This command is not available while LDAP authentication is active.

Renaming user groups

Use the **user_group_rename** command to rename a user group.

```
user_group_rename user_group=UserGroup new_name=Name
```

Parameters

Name	Type	Description	Mandatory
user_group	Object name	User group to be renamed.	Y
new_name	Object name	New name of the user group.	Y

Example:

```
user_group_rename user_group=ug1 new_name=ug2
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- USER_GROUP_NAME_DOES_NOT_EXIST
  The user group name does not exist.
- USER_GROUP_NAME_ALREADY_EXISTS
  The user group name already exists.

Updating a user group

Use the **user_group_update** command to update a user group.

```
user_group_update  user_group=UserGroup [ ldap_role=LdapRole ] [ access_all=<yes|no> ] [ domain=DomainList ]
```

Parameters

Name	Type	Description	Mandatory	Default
user_group	Object name	The name of the user group to be updated.	Y	N/A
ldap_role	String	The value representing the user group in LDAP.	N	Keep current LDAP role.
access_all	Boolean	Assigns application administration access level for all volumes.	N	no
domain	N/A	The user_group will be attached to the specified domains. To specify more than one domain, separate them with a comma. To specify all the existing domains, use "*".	N	none

Example:

```
user_group_update user_group=ug1 ldap_role="App Admin 1" access_all=yes
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- USER_GROUP_NAME_DOES_NOT_EXIST
  The user group name does not exist.
- LDAP_ROLE_ALREADY_USED
  This LDAP role is already in use in the LDAP configuration or in a user group.

- DOMAIN_DOESNT_EXIST
  The domain does not exist.
- REMOVAL_WOULD_CREATE_UNRESOLVABLE_REFERENCE_BETWEEN_HOST_AND_USERGROUP
  Completing this operation will result in a host referring to a user group that is not in its domain.

  Troubleshooting: Remove the reference explicitly and re-run the command.
- REMOVAL_WOULD_CREATE_UNRESOLVABLE_REFERENCE_BETWEEN_CLUSTER_AND_USERGROUP
  Completing this operation will result in a cluster referring to a user group that is not in its domain.

  Troubleshooting: Remove the reference explicitly and re-run the command.
- REMOVAL_WOULD_CREATE_UNRESOLVABLE_REFERENCE_BETWEEN_USER_AND_USERGROUP
  Completing this operation will result in a user referring to a user group that is not in its domain.

  Troubleshooting: Remove the reference explicitly and re-run the command.

Listing users

Use the **user_list** command to list all users or a specific user.

```
user_list [ user=UserName | show_users=<all|active> ] [ domain=DomainName ]
```

Parameters

Name	Type	Description	Mandatory	Default
user	Object name	The user to be listed.	N	All users.
show_users	Enumeration	Indicates whether all internal users will be listed, or only internal users that are active.	N	active
domain	Object name	The domain name.	N	All Domains

The following information is listed:

- User name: Lower case
- Category
- Email address
- Phone number
- Phone area code
- Containing user group

Passwords are not shown in the list.

Field ID	Field output	Default position
name	Name	1
category	Category	2
group	Group	3
active	Active	4
email_address	Email Address	5
area_code	Area Code	6
number	Phone Number	7
access_all	Access All	8
id	ID	N/A
creator	Creator	N/A
creator_category	Creator Category	N/A

Example:

```
user_list
```

Output:

```
Name                Category            Group
xiv_development     xiv_development           yes
xiv_maintenance     xiv_maintenance           yes
admin               storageadmin              yes
technician          technician                yes
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Allowed
Read-only users	Allowed

User Category	Permission
Technicians	Disallowed

Renaming users

Use the **user_rename** command to rename a user.

```
user_rename user=UserName new_name=Name
```

Parameters

Name	Type	Description	Mandatory
user	Object name	The user to be renamed. User names are lowercase.	Y
new_name	Object name	New name of the user.	Y

This command renames a user.

Example:

```
user_rename user=admin new_name=storage_admin
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- USER_NAME_DOES_NOT_EXIST
  The user name does not exist.
- USER_NAME_ALREADY_EXISTS
  The user name already exists.
- USER_CANNOT_BE_RENAMED
  The user cannot be renamed.
- LDAP_AUTHENTICATION_IS_ACTIVE
  This command is not available while LDAP authentication is active.
- USER_OWNS_RECOVERY_KEY
  The user owns a recovery key and therefore cannot be deleted or renamed.
- OPERATION_NOT_ALLOWED_ON_DESIGNATED_MSM_USER
  The designated MSM user cannot be renamed and cannot be exclusively associated with a domain.

Updating a user definition

Use the **user_update** command to update a user definition.

```
user_update user=UserName [ password=Password password_verify=Password ] [ email_address=email ]
 [ area_code=AreaCode ] [ number=PhoneNumber ] [ exclusive=<yes|no> ]
```

Parameters

Name	Type	Description	Mandatory	Default
user	Object name	The name of the user to be updated. User names are lower case.	Y	N/A

Name	Type	Description	Mandatory	Default
password	N/A	New password. Users can only change their own passwords. Password of the user to be updated. The password must have between 6 and 16 characters. Any symbols are allowed, except the following: • double quotation (") • single quotation or apostrophe (') • grave accent (`) Passwords are case sensitive.	N	Retains the current password.
password_verify	N/A	Verification of the password: Must be equal to the **password**.	N	Retains the current password.
email_address	N/A	Email address of the user (for event notification).	N	Leaves the current email address.
number	N/A	Cellular phone number of the user (for event notification via SMS) excluding the area code.	N	Leaves the current number.
area_code	N/A	Area code of the cellular phone number of the user.	N	Leaves the current area code.
exclusive	Boolean	This parameter can be set only by security administrator. If set to "yes", the user will be removed from the global domain. If set to "no", the user will get permissions on the global domain.	N	Leaves the current value.

Users can change their own user parameters, such as password, email address, area code, and phone number. **Only** Storage administrators and Security administrators can change parameters of other users. In addition, a security administrator can change the **exclusive** parameter for all users.

When updating a user password, the following restrictions apply:

User type	Predefined name (if exists)	Can change own password?	Can change other user passwords?	Limitations
Storage administrator	admin	Y	Y	Cannot change the password of the predefined Technician user
Storage administrator	(customer defined)	Y	N	N/A
Security administrator	(customer defined)	Y	Y	Cannot change the password of: • other Security administrators • the predefined Technician user • the predefined Storage administrator 'admin' user
Technician	technician	N	N	N/A
Other (customer defined)	(customer defined)	Y	N	N/A

Example:

```
user_update user=admin password=Passw0rd password_verify=Passw0rd
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission	Condition
Storage administrator	Conditionally Allowed	A user other than admin may only change its own configuration.
Storage integration administrator	Disallowed	N/A
Application administrator	Conditionally Allowed	A user of this category may only change its own configuration.
Security administrator	Conditionally Allowed	A user of this category may only change its own configuration.
Read-only users	Conditionally Allowed	A user other than admin may only change its own configuration.
Technicians	Disallowed	N/A

Return codes

- USER_NAME_DOES_NOT_EXIST
  The user name does not exist.
- PASSWORDS_DO_NOT_MATCH
  Make sure that passwords are identical.
- USER_PHONE_NUMBER_MUST_ACCOMPANY_AREA_CODE
  The phone number must be indicated together with the area code.
- ADMIN_CAN_NOT_CHANGE_TECHNICIAN_USER

The administrator is not allowed to modify the details of a technician.

- SMS_DESTINATION_REFERS_TO_USER
  An SMS destination refers to the user and therefore must be defined by a phone number and an area code.
- EMAIL_DESTINATION_REFERS_TO_USER
  An email destination refers to the user, and therefore must be defined by an email address.
- USER_NOT_ALLOWED_TO_CHANGE_OTHER_USERS
  This user is not allowed to modify the details of other users.
- USER_NOT_ALLOWED_TO_HAVE_PHONE_NUMBER
  This user is not allowed to have a phone number.
- USER_NOT_ALLOWED_TO_HAVE_EMAIL_ADDRESS
  This user is not allowed to have an email address.
- USER_NOT_ALLOWED_TO_CHANGE_PASSWORDS
  This user cannot change passwords of other users.
- USER_CANNOT_BE_UPDATED_WHILE_LDAP_AUTHENTICATION_IS_ACTIVE
  The user cannot be updated while LDAP authentication is active.
- NON_EXCLUSIVE_USER_NOT_SECURITY_ADMIN
  Only a security administrator can add a user to a domain non-exclusively.
- PREDEFINED_USER_CANNOT_BE_RESTRICTED_TO_DOMAIN
  The specified user cannot be exclusively associated with a domain.
- SIA_MUST_BE_ASSOCIATED_WITH_A_DOMAIN
  The storage integration administrator must be associated with a domain.
- USER_IS_REFERRED_TO_BY_DEST
  The user is referred to by an event destination and therefore cannot be deleted.
- OPERATION_NOT_ALLOWED_ON_DESIGNATED_MSM_USER
  The designated MSM user cannot be renamed and cannot be exclusively associated with a domain.
- REMOVAL_WOULD_CREATE_UNRESOLVABLE_REFERENCE_BETWEEN_USER_AND_USERGROUP
  Completing this operation will result in a user referring to a user group that is not in its domain.

  Troubleshooting: Remove the reference explicitly and re-run the command.

Creating a new domain

Use the **domain_create** command to create a domain.

```
domain_create domain=DomainName [ size=GB ] [ max_pools=MaxPools ] [ max_volumes=MaxVolumes ]
 [ max_cgs=MaxCGs ] [ max_mirrors=MaxMirrors ] [ max_dms=MaxDataMigrations ]
 [ perf_class=perfClassName ] [ ldap_id=LdapRole ]
```

Parameters

Name	Type	Description	Mandatory	Default
domain	Object name	The name of the domain to be created.	Y	N/A
size	Integer	Defines the sum of the sizes of all the pools associated with the domain, in gigabytes.	N	0
max_pools	Positive integer	The maximum number of pools that can be associated with this domain.	N	0
max_volumes	Positive integer	The maximum number of volumes that can be associated with all the pools in this domain.	N	0
max_cgs	Integer	The maximum number of consistency groups that can be associated with this domain.	N	512
max_mirrors	Positive integer	The maximum number of mirrors that can be associated with this domain.	N	0
max_dms	Positive integer	The maximum number of data migrations that can be associated with this domain.	N	0
perf_class	Object name	Name of a performance class.	N	none
ldap_id	String	The name to be associated with this domain in LDAP.	N	The domain name

Example:

```
domain_create domain=d1 size=1000
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- DO_YOU_WANT_TO_SHARE_LDAP_ID
  The specified LDAP ID is already in use. Are you sure you want to share the same LDAP ID?
- DOMAIN_SIZE_SMALL
  The domain size is smaller than the minimal pool size. Are you sure?

Return codes

- DOMAIN_ALREADY_EXISTS
  A domain with this name already exists.
- DOMAIN_MAX_REACHED
  The maximum allowed number of domain objects is already reached.
- PERF_CLASS_BAD_NAME
  The performance class does not exist.
- PERF_CLASS_ASSOCIATED_WITH_HOSTS
  Performance class *Performance Class* is already being used by a host.
- PERF_CLASS_ASSOCIATED_WITH_VOLUMES
  Performance class *Performance Class* is already being used by a volume.
- DOMAIN_INSUFFICIENT_CAPACITY
  There is not enough capacity available for the domain.
- DOMAIN_INSUFFICIENT_VOLUMES
  There are not enough volumes available for the domain.
- DOMAIN_INSUFFICIENT_POOLS
  There are not enough pools available for the domain.
- DOMAIN_INSUFFICIENT_CGS
  There are not enough consistency groups available for the domain.
- DOMAIN_INSUFFICIENT_MIRRORS
  There are not enough mirrors available for the domain.
- DOMAIN_INSUFFICIENT_DMS
  There are not enough data migrations available for the domain.

Updating a domain definition

Use the **domain_update** command to update a domain definition.

```
domain_update domain=DomainName [ size=GB ] [ max_pools=MaxPools ] [ max_volumes=MaxVolumes ]
 [ max_cgs=MaxCGs ] [ max_mirrors=MaxMirrors ] [ max_dms=MaxDataMigrations ]
 [ perf_class=perfClassName ] [ ldap_id=LdapRole ]
```

Parameters

Name	Type	Description	Mandatory	Default
domain	Object name	The name of the domain to be updated.	Y	N/A
size	Integer	Defines the sum of the sizes of all the pools associated with the domain, in gigabytes.	N	Current value.
max_pools	Positive integer	The maximum number of pools that can be associated with this domain.	N	Current value.
max_volumes	Positive integer	The maximum number of volumes that can be associated with all the pools in this domain.	N	Current value.
max_cgs	Integer	The maximum number of consistency groups that can be associated with this domain.	N	Current value.
max_mirrors	Positive integer	The maximum number of mirrors that can be associated with this domain.	N	Current value.
max_dms	Positive integer	The maximum number of data migrations that can be associated with this domain.	N	Current value.
perf_class	Object name	Name of a performance class.	N	Current value.
ldap_id	String	The name to be associated with this domain in LDAP.	N	Current value.

Example:

```
domain_update domain=d1 size=10000 max_pools=5 max_volumes=100
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- DO_YOU_WANT_TO_SHARE_LDAP_ID
  The specified LDAP ID is already in use. Are you sure you want to share the same LDAP ID?
- DOMAIN_SIZE_SMALL
  The domain size is smaller than the minimal pool size. Are you sure?

Return codes

- DOMAIN_DOESNT_EXIST
  The domain does not exist.
- PERF_CLASS_BAD_NAME
  The performance class does not exist.
- PERF_CLASS_ASSOCIATED_WITH_HOSTS
  Performance class *Performance Class* is already being used by a host.
- PERF_CLASS_ASSOCIATED_WITH_VOLUMES
  Performance class *Performance Class* is already being used by a volume.
- DOMAIN_SIZE_TOO_SMALL
  The domain usage exceeds the requested size.
- DOMAIN_MAX_VOLUMES_TOO_SMALL
  The actual number of volumes in the domain already exceeds the specified number.
- DOMAIN_MAX_MIRRORS_TOO_SMALL
  The actual number of mirrors in the domain already exceeds the specified number.
- DOMAIN_MAX_DMS_TOO_SMALL
  The actual number of data migrations in the domain already exceeds the specified number.
- DOMAIN_MAX_CGS_TOO_SMALL
  The actual number of consistency groups in the domain already exceeds the specified number.
- DOMAIN_MAX_POOLS_TOO_SMALL
  The actual number of pools in the domain already exceeds the specified number.
- DOMAIN_INSUFFICIENT_CAPACITY
  There is not enough capacity available for the domain.
- DOMAIN_INSUFFICIENT_VOLUMES
  There are not enough volumes available for the domain.
- DOMAIN_INSUFFICIENT_POOLS
  There are not enough pools available for the domain.
- DOMAIN_INSUFFICIENT_MIRRORS
  There are not enough mirrors available for the domain.
- DOMAIN_INSUFFICIENT_CGS
  There are not enough consistency groups available for the domain.
- DOMAIN_INSUFFICIENT_DMS
  There are not enough data migrations available for the domain.

Renaming a domain

Use the **domain_rename** command to rename a domain.

```
domain_rename domain=DomainName new_name=Name
```

Parameters

Name	Type	Description	Mandatory
new_name	Object name	Name of the domain.	Y
domain	Object name	New name of the domain.	Y

Example:

```
domain_rename domain=domain1 new_name=domain2
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- DOMAIN_ALREADY_EXISTS
  A domain with this name already exists.
- DOMAIN_DOESNT_EXIST
  The domain does not exist.

Deleting a domain

Use the **domain_delete** command to delete a domain.

```
domain_delete domain=DomainName
```

Parameters

Name	Type	Description	Mandatory
domain	Object name	The name of the domain to delete.	Y

Example:

```
domain_delete domain=domain1
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_DELETE_DOMAIN
  Are you sure you want to delete the domain *Domain*?

Return codes

- DOMAIN_DOESNT_EXIST
  The domain does not exist.
- DOMAIN_HAS_POOL
  One or more pools are still in the domain.
- DOMAIN_HAS_USER_GROUP
  One or more user groups are associated with this domain.
- DOMAIN_HAS_USER
  One or more users are associated with this domain.
- DOMAIN_HAS_SCHEDULE
  One or more schedules are associated with this domain.
- DOMAIN_HAS_DEST
  One or more destinations are associated with this domain.
- DOMAIN_HAS_DESTGROUP
  One or more destination groups are associated with this domain.

Listing domains

List all domains or the specified one.

```
domain_list [ domain=DomainName ]
```

Parameters

Name	Type	Description	Mandatory	Default
domain	Object name	Name of a domain.	N	All domains.

When the domain parameter is provided, only the specified domain is listed.

Example:

```
domain_list domain=d1
```

Output:

```
Tabular output

Name        DN          Soft    Free Soft    Hard    Free Hard
Domain1     Domain1     1703    0            1703    0
Domain2     Domain2     1703    1703         1703    1703
Domain3     Domain3     1703    1600         1703    1600
Domain4     Domain4     1703    1703         1703    1703
Domain5     Domain5     1703    1703         1703    1703

XML output
<domain
id="_home_markdown_jenkins_workspace_Transform_in_STJKN5_12.3.2_CLI_latest_r_command__accesscontrol__domain_list_4e414e00
000">
    <id value="4e414e00000"/>
    <name value="Domain1"/>
    <hard_capacity value="1703"/>
    <soft_capacity value="1703"/>
    <free_soft_capacity value="0"/>
    <free_hard_capacity value="0"/>
    <max_pools value="25"/>
    <used_pools value="1"/>
    <max_volumes value="100"/>
    <used_volumes value="2"/>
    <max_cgs value="100"/>
    <used_cgs value="1"/>
    <max_sync_mirrors value="70"/>
    <used_sync_mirrors value="0"/>
    <ax_async_mirrors value="70"/>
    <used_async_mirrors value="0"/>
    <perf_class_uid value="50713d00000"/>
    <perf_class value="QoS1"/>>
    <dn value="Domain1"/>
</domain>
```

Field ID	Field output	Default position
name	Name	1
ldap_id	LDAP ID	2
size	Size	3
size_MiB	Size (MiB)	N/A
total_pool_size	Total Pools (GB)	4

Field ID	Field output	Default position
total_pool_size_MiB	Total Pools (MiB)	N/A
empty_space	Empty (GB)	5
empty_space_MiB	Empty (MiB)	N/A
max_pools	Max Pools	6
used_pools	Pools	7
max_volumes	Max Volumes	8
used_volumes	Volumes	9
max_mirrors	Max Mirrors	10
used_mirrors	Mirrors	11
max_dms	Max Data Migrations	12
used_dms	Data Migrations	13
max_cgs	Max CGs	14
used_cgs	CGs	15
perf_class	Performance Class	16
managed	Managed	17
id	ID	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Allowed
Read-only users	Allowed
Technicians	Disallowed

Listing users per domain

Use the **domain_list_users** command to list users associated with domain(s).

```
domain_list_users  [ domain=DomainName ] [ user=UserName ] [ category=Category ] [ show_users=<all|active> ]
```

Parameters

Name	Type	Description	Mandatory	Default
domain	Object name	Name of a domain.	N	All domains.
user	Object name	Name of a user.	N	All users.
category	Enumeration	The roles of the users to be listed. Available options are: storageadmin, readonly, applicationadmin and storageintegrationadmin.	N	All categories.
show_users	Enumeration	Indicates whether to list all internal users, or only active internal users.	N	active

Example:

```
domain_list_users domain=d1
```

Output:

```
Domain   User         Category
-------- ---------- --------------
d1       d1_domain   storageadmin
```

Field ID	Field output	Default position
domain_name	Domain	1
user_name	User	2
category	Category	3

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Allowed
Read-only users	Allowed

User Category	Permission
Technicians	Allowed

Return codes

- DOMAIN_DOESNT_EXIST
  The domain does not exist.

Listing objects in domains

Use the **domain_list_objects** command to list objects attached to domain(s).

```
domain_list_objects  [ domain=DomainName ] [ type=ObjectType [ name=ObjectName ] ]
```

Parameters

Name	Type	Description	Mandatory	Default
domain	Object name	Name of a domain.	N	All domains.
type	Enumeration	The object type to list: target, host, cluster, schedule, usergroup, dest, destgroup or rule.	N	All object types.
name	Object name	Name of an object.	N	All object names.

This command is used for listing objects in the system per domain.

Example:

```
domain_list_objects domain=d1
```

Output:

```
Domain   Type        Object
-------- ----------- --------------
d1       cluster     c1
d1       host        MyHost
d1       schedule    min_interval
d1       schedule    never
```

Field ID	Field output	Default position
domain_name	Domain	1
object_type	Type	2
object_name	Object	3

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Allowed
Read-only users	Allowed
Technicians	Allowed

Return codes

- DOMAIN_DOESNT_EXIST
  The domain does not exist.
- TOO_MANY_OBJECTS
  There are too many objects to output. Re-run the command by using the command's parameters to filter the output by domain or object type.

Listing the global domain

Use the **domain_global_list** to list the global domain.

```
domain_global_list
```

Example:

```
domain_global_list
```

Field ID	Field output	Default position
name	Name	1
ldap_id	LDAP ID	2
size	Size	3
size_MiB	Size (MiB)	N/A
total_pool_size	Total Pools (GB)	4
total_pool_size_MiB	Total Pools (MiB)	N/A
empty_space	Empty (GB)	5
empty_space_MiB	Empty (MiB)	N/A
max_pools	Max Pools	6
used_pools	Pools	7
max_volumes	Max Volumes	8
used_volumes	Volumes	9
max_mirrors	Max Mirrors	10
used_mirrors	Mirrors	11
max_dms	Max Data Migrations	12
used_dms	Data Migrations	13
max_cgs	Max CGs	14
used_cgs	CGs	15
perf_class	Performance Class	16
managed	Managed	17
id	ID	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Attaching an object to a domain

Use the **domain_attach_object** command to associate an object with a domain.

```
domain_attach_object domain=DomainName type=ObjectType name=ObjectName
```

Parameters

Name	Type	Description	Mandatory
domain	Object name	The name of the domain.	Y
type	Enumeration	The object type to attach to the domain. It can be: target, host, cluster, schedule, usergroup, dest, destgroup or rule.	Y
name	Object name	The object name.	Y

Example:

```
domain_attach_object domain=d1 type=host name=MyHost
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- DOMAIN_DOESNT_EXIST
  The domain does not exist.
- TARGET_BAD_NAME
  The target name does not exist.
- HOST_BAD_NAME
  The host name does not exist.
- CLUSTER_BAD_NAME
  The cluster name does not exist.
- USER_GROUP_NAME_DOES_NOT_EXIST
  The user group name does not exist.
- SCHEDULE_DOES_NOT_EXIST
  The specified schedule does not exist.
- DEST_NAME_DOES_NOT_EXIST
  The destination name does not exist.
- DESTGROUP_NAME_DOES_NOT_EXIST
  The destination group name does not exist.
- EVENT_RULE_NAME_DOES_NOT_EXIST
  The event rule name does not exist.
- USER_IS_NOT_IN_DESTINATION_DOMAINS
  The user must be included in destination domains.
- DESTINATION_IS_NOT_IN_DESTGROUP_DOMAINS
  The destination must be included in the destination group domains.
- DESTINATION_IS_NOT_IN_RULE_DOMAINS
  The destination must be included in rule domains.
- DESTGROUP_IS_NOT_IN_RULE_DOMAINS
  The destination groups must be included in rule domains.
- ESCALATION_RULE_NOT_IN_RULE_DOMAINS
  An escalation rule must belong to rule domains.
- CLUSTER_HAS_HOSTS_UNASSOCIATED_WITH_DOMAIN
  The cluster cannot be attached, because it includes hosts that are not in the specified domain.
- RESOURCE_ALREADY_ASSOCIATED_WITH_THIS_DOMAIN
  The resource is already associated with this domain.

Disassociating object from a domain

Use the **domain_detach_object** command to disassociate object from a domain.

```
domain_detach_object domain=DomainName type=ObjectType name=ObjectName
```

Parameters

Name	Type	Description	Mandatory
domain	Object name	The name of the domain.	Y
type	Enumeration	The object type to disassociate from the domain. It can be: target, host, cluster, schedule, usergroup, dest, destgroup, or rule.	Y
name	Object name	The object name.	Y

The object is disassociated from mapped or bound objects that belong to the domain.

Example:

```
domain_detach_object domain=d1 type=host name=MyHost
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed

User Category	Permission
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- DOMAIN_DOESNT_EXIST
  The domain does not exist.
- RESOURCE_NOT_ASSOCIATED_WITH_THIS_DOMAIN
  The resource is not associated with this domain.
- DOMAIN_VOLUME_MAPPED_TO_HOST
  A domain volume is mapped to this host.
- DOMAIN_VOLUME_MAPPED_TO_CLUSTER
  Cluster has a volume in the domain mapped to it.
- REMOVAL_WOULD_CREATE_UNRESOLVABLE_REFERENCE_BETWEEN_HOST_AND_USERGROUP
  Completing this operation will result in a host referring to a user group that is not in its domain.

  Troubleshooting: Remove the reference explicitly and re-run the command.
- REMOVAL_WOULD_CREATE_UNRESOLVABLE_REFERENCE_BETWEEN_CLUSTER_AND_USERGROUP
  Completing this operation will result in a cluster referring to a user group that is not in its domain.

  Troubleshooting: Remove the reference explicitly and re-run the command.
- HOST_PART_OF_ATTACHED_CLUSTER
  The host is part of a cluster and cannot be handled individually.
- DOMAIN_TARGET_IN_USE
  The target domain cannot be removed, because it includes a volume used in a mirror, data migration, or OLVM relationship.
- DOMAIN_VOLUME_BOUND_TO_HOST
  A domain volume is bound to this host via an ALU.
- TARGET_BAD_NAME
  The target name does not exist.
- HOST_BAD_NAME
  The host name does not exist.
- CLUSTER_BAD_NAME
  The cluster name does not exist.
- USER_GROUP_NAME_DOES_NOT_EXIST
  The user group name does not exist.
- SCHEDULE_DOES_NOT_EXIST
  The specified schedule does not exist.
- DEST_NAME_DOES_NOT_EXIST
  The destination name does not exist.
- DESTGROUP_NAME_DOES_NOT_EXIST
  The destination group name does not exist.
- EVENT_RULE_NAME_DOES_NOT_EXIST
  The event rule name does not exist.
- DETACH_WOULD_MAKE_OBJECT_INACCESSIBLE
  Detaching the object will render it unassociated with any domain, and therefore inaccessible.

  Troubleshooting: Delete the object, if it is no longer needed.
- USER_IS_NOT_IN_DESTINATION_DOMAINS
  The user must be included in destination domains.
- DESTINATION_IS_NOT_IN_RULE_DOMAINS
  The destination must be included in rule domains.
- DESTINATION_IS_NOT_IN_DESTGROUP_DOMAINS
  The destination must be included in the destination group domains.
- REMOVAL_WOULD_CREATE_UNRESOLVABLE_REFERENCE_BETWEEN_USER_AND_USERGROUP
  Completing this operation will result in a user referring to a user group that is not in its domain.

  Troubleshooting: Remove the reference explicitly and re-run the command.
- DESTGROUP_IS_NOT_IN_RULE_DOMAINS

The destination groups must be included in rule domains.

- ESCALATION_RULE_NOT_IN_RULE_DOMAINS
  An escalation rule must belong to rule domains.
- DOMAIN_SCHEDULE_IN_USE
  The schedule is in use, and therefore cannot be moved to another domain.
- DOMAIN_PROXY_VOLUME_MAPPED_TO_HOST
  A proxy domain volume is mapped to this host.
- DOMAIN_PROXY_VOLUME_MAPPED_TO_CLUSTER
  A proxy domain volume is mapped to this cluster.

Associating users to a domain

Use the **domain_add_user** command to associate a user to a domain.

```
domain_add_user domain=DomainName user=UserName [ exclusive=<yes|no> ]
```

Parameters

Name	Type	Description	Mandatory	Default
domain	Object name	The name of the domain.	Y	N/A
user	Object name	The name of the user.	Y	N/A
exclusive	Boolean	Set to Yes to restrict the user to domain's objects.	N	yes

Example:

```
domain_add_user domain=d1 user=d1_admin
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- LDAP_AUTHENTICATION_IS_ACTIVE
  This command is not available while LDAP authentication is active.
- DOMAIN_DOESNT_EXIST
  The domain does not exist.
- USER_NAME_DOES_NOT_EXIST
  The user name does not exist.
- NON_EXCLUSIVE_USER_NOT_SECURITY_ADMIN
  Only a security administrator can add a user to a domain non-exclusively.
- USER_CANNOT_BE_ADDED_TO_A_DOMAIN
  The specified user cannot be associated with a domain.
- PREDEFINED_USER_CANNOT_BE_RESTRICTED_TO_DOMAIN
  The specified user cannot be exclusively associated with a domain.
- DOMAIN_USER_EXIST
  This user is already added to the domain.
- USER_IS_REFERRED_TO_BY_DEST
  The user is referred to by an event destination and therefore cannot be deleted.
- OPERATION_NOT_ALLOWED_ON_DESIGNATED_MSM_USER
  The designated MSM user cannot be renamed and cannot be exclusively associated with a domain.
- SIA_MUST_BE_ASSOCIATED_WITH_A_DOMAIN
  The storage integration administrator must be associated with a domain.

- REMOVAL_WOULD_CREATE_UNRESOLVABLE_REFERENCE_BETWEEN_USER_AND_USERGROUP
  Completing this operation will result in a user referring to a user group that is not in its domain.

  Troubleshooting: Remove the reference explicitly and re-run the command.

Removing a user from a domain

Use the **domain_remove_user** command to remove a user from a domain.

```
domain_remove_user domain=DomainName user=UserName
```

Parameters

Name	Type	Description	Mandatory
domain	Object name	The name of the domain.	Y
user	Object name	The name of the user.	Y

Example:

```
domain_remove_user domain=d1 user=d1_admin
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- USER_WILL_BE_DELETED
  The user is associated with this domain only. This will delete the user from the system. Are you sure?

Return codes

- USER_NAME_DOES_NOT_EXIST
  The user name does not exist.
- DOMAIN_DOESNT_HAVE_THE_USER
  The user is not attached to this domain.
- DOMAIN_DOESNT_EXIST
  The domain does not exist.
- LDAP_AUTHENTICATION_IS_ACTIVE
  This command is not available while LDAP authentication is active.
- REMOVAL_WOULD_CREATE_UNRESOLVABLE_REFERENCE_BETWEEN_USER_AND_USERGROUP
  Completing this operation will result in a user referring to a user group that is not in its domain.

  Troubleshooting: Remove the reference explicitly and re-run the command.
- DOMAIN_USER_CANNOT_REMOVE_HIMSELF
  Users cannot remove themselves from a domain.
- USER_IS_REFERRED_TO_BY_DEST
  The user is referred to by an event destination and therefore cannot be deleted.

Adding a pool to a domain

Use the **domain_add_pool** command to add a pool to a domain.

```
domain_add_pool domain=DomainName pool=PoolName [ adjust=<yes|no> ]
```

Parameters

Name	Type	Description	Mandatory	Default
domain	Object name	The name of the domain.	Y	N/A
pool	Object name	The pool name.	Y	N/A
adjust	Boolean	Adjust domain resources. If 'adjust' is set to true, the resources of the global domain and destination domain are adjusted to accommodate the pool being moved.	N	no

Example:

```
domain_add_pool domain=d1 pool=p1
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- POOL_DOES_NOT_EXIST
  The storage pool does not exist.
- NO_FREE_CAPACITY_IN_DOMAIN
  There is not enough free space in the domain.
- DOMAIN_DOESNT_EXIST
  The domain does not exist.
- POOL_ALREADY_ASSOCIATED_WITH_A_DOMAIN
  The pool is already associated with a domain.
- DOMAIN_MAX_POOLS_REACHED
  The maximum allowed number of domain pools is already reached.
- DOMAIN_MAX_MIRRORS_REACHED
  The domain exceeds the maximum allowed number of mirrors.
- DOMAIN_MAX_DMS_REACHED
  The domain exceeds the maximum allowed number of data migrations.
- DOMAIN_MAX_CONS_GROUPS_REACHED
  The domain exceeds the maximum allowed number of consistency groups.
- DOMAIN_MAX_VOLUMES_REACHED
  The domain exceeds the maximum allowed number of volumes.
- DOMAIN_USED_TARGET_NOT_IN_DESTINATION
  A target that is used by mirror in the pool is not associated with the target domain.
- DOMAIN_USED_SCHEDULE_NOT_IN_DESTINATION
  A schedule used by a mirror in the pool is not associated with the target domain.
- MAPPED_HOSTS_NOT_IN_DESTINATION
  A host that is mapped to a volume in the pool is not associated with the target domain.
- MAPPED_CLUSTERS_NOT_IN_DESTINATION
  A cluster that is mapped to a volume in the pool is not associated with the target domain.
- NO_SPACE
  The system does not have enough free space for the requested storage pool size.

Removing a pool from a domain

Use the **domain_remove_pool** command to remove a pool from a domain.

```
domain_remove_pool domain=DomainName pool=PoolName [ adjust=<yes|no> ]
```

Parameters

Name	Type	Description	Mandatory	Default
domain	Object name	The name of the domain.	Y	N/A
pool	Object name	The pool name.	Y	N/A
adjust	Boolean	Adjust domain resources. If set to True, the resources of the global domain and destination domain are adjusted to accommodate the pool being moved.	N	no

Example:

```
domain_remove_pool domain=d1 pool=p1
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- DOMAIN_DOESNT_EXIST
  The domain does not exist.
- POOL_DOES_NOT_EXIST
  The storage pool does not exist.
- DOMAIN_DOESNT_HAVE_THE_POOL
  The pool is not attached to this domain.
- MAX_POOLS_REACHED
  The maximum allowed number of storage pools is already reached.
- MAX_MIRRORS_REACHED
  The maximum number of mirrors is already reached.
- MAX_CONS_GROUPS_REACHED
  The maximum allowed number of consistency groups is already reached.
- MAX_VOLUMES_REACHED
  The maximum allowed number of volumes is already reached.
- MAX_DMS_REACHED
  The maximum number of remote volumes (mirror/migration) is already reached.

  Troubleshooting: Delete unnecessary data migration objects.
- NO_SPACE
  The system does not have enough free space for the requested storage pool size.
- DOMAIN_USED_SCHEDULE_NOT_IN_DESTINATION
  A schedule used by a mirror in the pool is not associated with the target domain.

Moving a pool from one domain to another

Use the **domain_move_pool** command to move a pool from one domain to another.

```
domain_move_pool pool=PoolName src_domain=DomainName dst_domain=DomainName [ adjust=<yes|no> ]
```

Parameters

Name	Type	Description	Mandatory	Default
pool	Object name	The name of the pool to be moved.	Y	N/A
src_domain	Object name	The source domain name.	Y	N/A
dst_domain	Object name	The destination domain name.	Y	N/A
adjust	Boolean	Adjust domain resources. If set to Yes, the resources of the domains are adjusted to accommodate the pool being moved.	N	no

Example:

```
domain_move_pool pool=p1 src_domain=d1 dst_domain=d2
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- SOURCE_DOMAIN_DOES_NOT_EXIST
  The source domain does not exist.
- DESTINATION_DOMAIN_DOES_NOT_EXIST
  The destination domain does not exist.
- POOL_DOES_NOT_EXIST
  The storage pool does not exist.
- POOL_NOT_ASSOCIATED_WITH_SOURCE_DOMAIN
  The pool is not associated with the source domain.
- DOMAIN_MAX_POOLS_REACHED
  The maximum allowed number of domain pools is already reached.
- MAPPED_HOSTS_NOT_IN_DESTINATION
  A host that is mapped to a volume in the pool is not associated with the target domain.
- MAPPED_CLUSTERS_NOT_IN_DESTINATION
  A cluster that is mapped to a volume in the pool is not associated with the target domain.
- NO_FREE_CAPACITY_IN_DOMAIN
  There is not enough free space in the domain.
- DOMAIN_USED_SCHEDULE_NOT_IN_DESTINATION
  A schedule used by a mirror in the pool is not associated with the target domain.
- DOMAIN_USED_TARGET_NOT_IN_DESTINATION
  A target that is used by mirror in the pool is not associated with the target domain.
- DOMAIN_MAX_MIRRORS_REACHED
  The domain exceeds the maximum allowed number of mirrors.
- DOMAIN_MAX_DMS_REACHED
  The domain exceeds the maximum allowed number of data migrations.
- DOMAIN_MAX_CONS_GROUPS_REACHED
  The domain exceeds the maximum allowed number of consistency groups.
- DOMAIN_MAX_VOLUMES_REACHED
  The domain exceeds the maximum allowed number of volumes.
- OPERATION_DENIED_OBJECT_MANAGED
  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.
- SOURCE_AND_DESTINATION_DOMAINS_MUST_BE_DIFFERENT
  The source and destination domains must be different.
- DOMAIN_CONTAINS_OLVM_VOLUME
  The domain contains a volume in the OLVM process.
- NO_SPACE
  The system does not have enough free space for the requested storage pool size.

Setting the domain attribute

Use the **domain_manage** command to set or clear the Managed attribute of a domain.

```
domain_manage domain=DomainName managed=<yes|no>
```

Parameters

Name	Type	Description	Mandatory
domain	Object name	The domain name.	Y
managed	Boolean	The Managed attribute. If set to Yes, the domain will be marked as managed.	Y

Example:

```
domain_mmanage domain=d1 managed=yes
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Disallowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- DOMAIN_DOESNT_EXIST
  The domain does not exist.
- DOMAIN_HAS_POOL
  One or more pools are still in the domain.
- DOMAIN_HAS_SCHEDULE
  One or more schedules are associated with this domain.
- DOMAIN_HAS_DEST
  One or more destinations are associated with this domain.
- DOMAIN_HAS_DESTGROUP
  One or more destination groups are associated with this domain.
- DOMAIN_HAS_USER_GROUP
  One or more user groups are associated with this domain.

Setting domain related policies

Use the **domain_policy_set** command to set domain-related policies.

```
domain_policy_set name=Name value=ParamValue
```

Parameters

Name	Type	Description	Mandatory
name	String	Name of the parameter to set.	Y
value	String	Value of the parameter.	Y

This command is used for setting domain related policies.

- name=access defines whether non-domain-specific users can access domain-specific resources (*value=open*) or not (*value=closed*).
- name=host_management defines whether domain administrators can create their own hosts (*value=extended*), or are restricted to hosts assigned to their domains by NDSO administrators (*value=basic*).

Example:

```
domain_policy_set name=access value=open
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission

User Category	Permission
Storage administrator	Disallowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- UNRECOGNIZED_CONFIG_PARAMETER
  Unrecognized configuration parameter: '*name*'.

  Troubleshooting: Use a valid configuration parameter in the command syntax. For the list of valid configuration parameters, see the CLI Reference Guide.

Displaying domain-related policies

Use the **domain_policy_get** command to display domain-related policies.

```
domain_policy_get [ name=Name ]
```

Parameters

Name	Type	Description	Mandatory	Default
name	String	Name of the parameter to get.	N	All parameters.

- name=access defines whether non-domain-specific users can access domain-specific resources (*value=open*) or not (*value=closed*).
- name=host_management defines whether domain administrators can create their own hosts (*value=extended*), or are restricted to hosts assigned to their domains by NDSO administrators (*value=basic*).

Field ID	Field output	Default position
name	Name	1
value	Value	2

Example:

```
domain_policy_get name=access
```

Output:

```
Name     Value
-------- -------
access   OPEN
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Allowed
Read-only users	Allowed
Technicians	Allowed

Return codes

- UNRECOGNIZED_CONFIG_PARAMETER
  Unrecognized configuration parameter: '*name*'.

  Troubleshooting: Use a valid configuration parameter in the command syntax. For the list of valid configuration parameters, see the CLI Reference Guide.

- CONF_SERVER_UNREACHABLE
  The configuration server is unreachable.

Specifying a user associated with IBM Hyper-Scale Manager

Use the **designate_msm_user_set** command to specify the name of the user that is associated with the IBM Hyper-Scale Manager.

```
designate_msm_user_set name=UserName
```

Parameters

Name	Type	Description	Mandatory
name	Object name	The designated user.	Y

This command specifies which XIV user is defined in the IBM Hyper-Scale Manager Server in the activation step. This can be either a local or LDAP user, depending on whether LDAP authentication is used.

Example:

```
designate_msm_user_set name=xiv_msms
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- USER_NAME_DOES_NOT_EXIST
  The user name does not exist.
- USER_IS_ONLY_DOMAIN_ADMIN
  The user is associated with one or more domains, and cannot view the entire system.

Retrieving the user associated with the IBM Hyper-Scale Manager

Use the **designate_msm_user_get** command to retrieve the name of the user associated with the IBM Hyper-Scale Manager.

```
designate_msm_user_get
```

Example:

```
designate_msm_user_get
```

Output:

```
xiv_msms
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Allowed
Read-only users	Allowed
Technicians	Allowed

Return codes

- NO_DESIGNATED_MSM_USER
  There is no designated IBM Hyper-Scale user.

Setting the application administrator's scope of commands

Use the **appadmin_capabilities_set** command to define whether an Application Administrator is authorized to perform the basic or advanced set of commands.

```
appadmin_capabilities_set value=<basic|advanced>
```

Parameters

Name	Type	Description	Mandatory
value	Enumeration	The set of commands that an Application Administrator is authorized to perform.	Y

These are the additional commands available when an application administrator is authorized to perform advanced commands:

- mirror_statistics_get
- mirror_activate
- mirror_deactivate
- mirror_change_role
- mirror_switch_roles

Example:

```
appadmin_capabilities_set value=basic
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Disallowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Getting the application administrator's scope of commands

Use the **appadmin_capabilities_get** command to display the state of the Application Administrator's capabilities.

```
appadmin_capabilities_get
```

Example:

```
appadmin_capabilities_get
```

Output:

```
BASIC
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Allowed
Read-only users	Allowed
Technicians	Allowed

Fibre Channel and iSCSI configuration and status commands

This section describes the command-line interface (CLI) for Fibre Channel port configuration.

- **Discovering FC hosts**
  Use the **fc_connectivity_list** command to discover FC hosts and targets on the FC network.
- **Changing FC port configuration**
  Use the **fc_port_config** command to configure FC ports.
- **Listing FC ports**
  Use the **fc_port_list** command to list the status and configuration of the system's FC ports.
- **Listing FC Port Tests**
  Use the **fc_port_test_list** command to list the status of the system's FC port tests.
- **Starting FC Port Test**
  Use **fc_port_test_start** to run a test of the FC port.
- **Aborting FC Port Test**
  Use **fc_port_test_abort** to abort a currently running FC port test.

- **Listing connectivity to hosts**
  Use the **host_connectivity_list** command to list Fibre Channel and iSCSI-level connectivity.

Discovering FC hosts

Use the **fc_connectivity_list** command to discover FC hosts and targets on the FC network.

```
fc_connectivity_list  [ role=<dual|initiator|target> ] [ wwpn=WWPN ]
 [ module=ModuleNumber | fc_port=ComponentId ]
```

Parameters

Name	Type	Description	Mandatory	Default
role	Enumeration	Specifies whether to discover initiators or targets.	N	List all - targets and/or initiators.
wwpn	N/A	Limits the output only to this specific address.	N	All addresses
module	N/A	Limits the output to the enabled connectivity to this module.	N	All modules
fc_port	N/A	Limits the output to this specific XIV port.	N	All ports

This command lists FC hosts on the network.

role=initiator detects initiators on the network. When **role=initiator**, the *non-logged-in* option can only be used to debug hosts that are on the network, but did not log in.

role=target detects targets. When **type=target**, the *non-logged-in* option can only be used to debug targets that rejected the storage system login. This command returns an error for an attempt to list targets from a target-only port, or to list initiators from an initiator-only port. Each output line contains the following information:

- Component ID (of the module)
- Storage system port number (within the module)
- WWPN
- Port ID (can be correlated with the switch database)
- Role: Initiator, Target, Dual
- Initiator/target (is the same for all lines of the same command)
- Login status (Yes/No)

Field ID	Field output	Default position
component_id	Component ID	1
wwpn	WWPN	2
port_id	Port ID	3
role	Role	4

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Changing FC port configuration

Use the **fc_port_config** command to configure FC ports.

```
fc_port_config fc_port=ComponentId [ enabled=<yes|no> ] [ role=<target|initiator> ] [ rate=<2|4|8|16|auto> ]
```

Parameters

Name	Type	Description	Mandatory	Default
fc_port	N/A	Port identifier.	Y	N/A
enabled	Boolean	Allows you to enable or disable the port.	N	yes
role	Enumeration	Port role: target, initiator or both.	N	Leaves the role unchanged.
rate	Enumeration	Line rate or auto for auto-negotiated rate.	N	Leaves the rate unchanged.

Example:

```
fc_port_config fc_port=1:FC_Port:1:1 enabled=yes role=Target rate=auto
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Allowed

Return codes

- COMPONENT_IS_NOT_AN_FC_PORT
  An FC port must be specified for the component.
- FC_PORT_DOES_NOT_EXIST
  The specified FC port does not exist.
- FC_PORT_TEST_IN_PROGRESS
  The FC port test is already in progress

Listing FC ports

Use the **fc_port_list** command to list the status and configuration of the system's FC ports.

```
fc_port_list [ module=ModuleNumber  | fcport=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
module	Limits the listing to a specific module.	N	All ports in all modules.
fcport	Lists only a specific port.	N	All ports in all modules.

This command lists all or some FC ports on the system. When no parameters are specified, all ports are listed. If a module is specified without a port, all ports on that module are listed. If a port is specified, a single port is listed.

The following information is provided for each port:

- Component ID of the module Port number (internal to module) 1-N
- WWPN
- Port ID
- Role (Initiator, Target, Dual)
- User-enabled (Yes/No)
- Maximum support rate: 2GB, 4GB, 8GB; constant - function of the HBA's capability
- Configured rate: 2GB, 4GB, 8GB, auto-negotiation; cannot be greater than the maximum supported rate
- Current active rate: 2GB, 4GB, 8GB; equal to the configured rate, unless the configured rate is auto-negotiation
- Port state: Online, Offline, Loopback, Link Down (physical connection is on, but no logical connection exists)
- Error counts
- Link type: Fabric Direct Attach, Private Loop, Point-to-Point, Public Loop, Unknown

Example:

```
fc_port_list
```

Output:

```
Component ID     Status   Currently Functioning   WWPN               Port ID    Role
---------------- -------- ----------------------- ------------------ ---------- --------
1:FC_Port:12:1   OK       yes                     5001738035C601C0   FFFFFFFF   Target
1:FC_Port:12:2   OK       yes                     5001738035C601C1   FFFFFFFF   Target
1:FC_Port:12:3   OK       yes                     5001738035C601C2   FFFFFFFF   Target
1:FC_Port:12:4   OK       yes                     5001738035C601C3   00EF009A   Target
1:FC_Port:13:1   OK       yes                     5001738035C601D0   FFFFFFFF   Target
1:FC_Port:13:2   OK       yes                     5001738035C601D1   FFFFFFFF   Target
1:FC_Port:13:3   OK       yes                     5001738035C601D2   FFFFFFFF   Target
1:FC_Port:13:4   OK       yes                     5001738035C601D3   FFFFFFFF   Target
1:FC_Port:8:1    OK       yes                     5001738035C60180   FFFFFFFF   Target
1:FC_Port:8:2    OK       yes                     5001738035C60181   FFFFFFFF   Target
1:FC_Port:8:3    OK       yes                     5001738035C60182   FFFFFFFF   Target
1:FC_Port:8:4    OK       yes                     5001738035C60183   00163AC0   Target
```

```
Cont.:

User Enabled   Current Rate (GBaud)   Port State     Link Type
```

```
-------------- ---------------------- -------------- ----------------------
yes            Auto                   Link Problem   None
yes            Auto                   Link Problem   None
yes            Auto                   Link Problem   None
yes            8                      Online         Fabric Direct Attach
yes            Auto                   Link Problem   None
yes            Auto                   Link Problem   None
yes            Auto                   Link Problem   None
yes            Auto                   Link Problem   None
yes            Auto                   Link Problem   None
yes            Auto                   Link Problem   None
yes            Auto                   Link Problem   None
yes            16                     Online         Fabric Direct Attach
```

```
Cont.:

Error Count   Active Firmware
------------- -----------------
0             8.3.40
0             8.3.40
0             8.3.40
0             8:3:40
0             8.3.40
0             8.3.40
0             8.3.40
0             8.3.40
0             8.3.40
0             8.3.40
0             8.3.40
0             8:3:40
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
port_num	Port Number	N/A
wwpn	WWPN	4
port_id	Port ID	5
role	Role	6
user_enabled	User Enabled	7
max_supported_rate	Maximum Supported Rate (GBaud)	N/A
configured_rate	Configured Rate (GBaud)	N/A
current_rate	Current Rate (GBaud)	8
port_state	Port State	9
link_type	Link Type	10
error_count	Error Count	11
active_firmware	Active Firmware	12
credit	Credit	N/A
hba_vendor	HBA Vendor	N/A
is_enabled	Enabled	N/A
module	Module	N/A
serial	Serial	N/A
temperature	Temperature	N/A
part_number	Part Number	N/A
original_serial	Original Serial	N/A
model	Model	N/A
original_model	Original Model	N/A
requires_service	Requires Service	N/A
service_reason	Service Reason	N/A
start_statstime	Timestamp of Link Statistics	N/A
link_failure	Link Failure	N/A
loss_of_sync	Loss of Sync	N/A
loss_of_signal	Loss of Signal	N/A
primit_seq_prot_error	Primitive Sequence Protocol Error	N/A
invalid_tx_word	Invalid Transmission Word	N/A
invalid_crc	Invalid CRC	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed

User Category	Permission
Read-only users	Allowed
Technicians	Allowed

Listing FC Port Tests

Use the **fc_port_test_list** command to list the status of the system's FC port tests.

```
fc_port_test_list [ fc_port=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
fc_port	Lists only a specific port.	N	All ports in all modules.

This command lists all or some FC port tests on the system. When no parameters are specified, all tests are listed. If a port is specified, a single test is listed.

Example:

```
fc_port_test_list
```

Output:

```
Component ID     Status       Start Time   Duration (Sec)   Transfer Length   CRC
---------------- ------------ ------------ ---------------- ----------------- -----
1:FC_Port:14:1   not active   NA           0                0                 0
1:FC_Port:14:2   not active   NA           0                0                 0
1:FC_Port:14:3   not active   NA           0                0                 0
1:FC_Port:14:4   not active   NA           0                0                 0
1:FC_Port:3:1    not active   NA           0                0                 0
1:FC_Port:3:2    not active   NA           0                0                 0
1:FC_Port:3:3    not active   NA           0                0                 0
1:FC_Port:3:4    not active   NA           0                0                 0
1:FC_Port:5:1    not active   NA           0                0                 0
1:FC_Port:5:2    not active   NA           0                0                 0
1:FC_Port:5:3    not active   NA           0                0                 0
1:FC_Port:5:4    not active   NA           0                0                 0
1:FC_Port:6:1    not active   NA           0                0                 0
1:FC_Port:6:2    not active   NA           0                0                 0
1:FC_Port:6:3    not active   NA           0                0                 0
1:FC_Port:6:4    not active   NA           0                0                 0
```

```
Cont.:
```

```
Frame Length Error   Disparity   Loopback Mode   Data Pattern   Data Size   Test Count
-------------------- ----------- --------------- -------------- ----------- ------------
0                    0           internal        0000           0           0
0                    0           internal        0000           0           0
0                    0           internal        0000           0           0
0                    0           internal        0000           0           0
0                    0           internal        0000           0           0
0                    0           internal        0000           0           0
0                    0           internal        0000           0           0
0                    0           internal        0000           0           0
0                    0           internal        0000           0           0
0                    0           internal        0000           0           0
0                    0           internal        0000           0           0
0                    0           internal        0000           0           0
0                    0           internal        0000           0           0
0                    0           internal        0000           0           0
0                    0           internal        0000           0           0
0                    0           internal        0000           0           0
```

```
Cont.:
```

```
Abort On Error
----------------
no
no
no
no
no
no
no
no
no
no
no
no
no
no
no
no
no
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
start_time	Start Time	3
duration	Duration (Sec)	4
transfer_length	Transfer Length	5
crc	CRC	6
frame_length_error	Frame Length Error	7
disparity	Disparity	8
loopback_mode	Loopback Mode	9
pattern	Data Pattern	10
data_size	Data Size	11
test_count	Test Count	12
increment	Test Increment	N/A
abort_on_error	Abort On Error	13

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Starting FC Port Test

Use **fc_port_test_start** to run a test of the FC port.

```
fc_port_test_start fc_port=ComponentId [ loopback=<internal|external> ] [ pattern=PATTERN ]
 [ data_size=DATA_SIZE ] [ frames=FRAMES ] [ increment=INCREMENT ]
 [ abort_on_error=<no|yes> ]
```

Parameters

Name	Type	Description	Mandatory	Default
fc_port	N/A	Port identifier.	Y	N/A
loopback	Enumeration	Loopback mode to use.	N	internal
pattern	String	Data pattern.	N	1414
data_size	Positive integer	Data size.	N	2048
frames	Positive integer	Number of frames.	N	10000
increment	Positive integer	Test increment.	N	1
abort_on_error	Boolean	Abort the test in case of an error.	N	no

Example:

```
fc_port_test_start fc_port=1:FC_Port:1:1
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Allowed

Return codes

- COMPONENT_IS_NOT_AN_FC_PORT
  An FC port must be specified for the component.

- FC_PORT_DOES_NOT_EXIST
  The specified FC port does not exist.
- COMPONENT_DOES_NOT_EXIST
  The component does not exist.
- TEST_NOT_ALLOWED_IN_CURRENT_STATUS
  This component cannot be tested in its current status.
- FC_PORT_TEST_IN_PROGRESS
  The FC port test is already in progress
- INVALID_FC_PORT_TEST_DATA_PATTERN
  The data pattern for the FC port test is invalid.

Aborting FC Port Test

Use **fc_port_test_abort** to abort a currently running FC port test.

```
fc_port_test_abort fc_port=ComponentId
```

Parameters

Name	Description	Mandatory
fc_port	Port identifier.	Y

Example:

```
fc_port_test_abort fc_port=1:FC_Port:1:1
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Allowed

Return codes

- COMPONENT_IS_NOT_AN_FC_PORT
  An FC port must be specified for the component.
- FC_PORT_DOES_NOT_EXIST
  The specified FC port does not exist.
- COMPONENT_DOES_NOT_EXIST
  The component does not exist.
- NO_FC_PORT_TEST_IN_PROGRESS
  The FC port test is currently not running.

Listing connectivity to hosts

Use the **host_connectivity_list** command to list Fibre Channel and iSCSI-level connectivity.

```
host_connectivity_list [ host=HostName | fc_host_port=WWPN ] [ module=ModuleNumber | fcport=ComponentId ] [
domain=DomainName ]
```

Parameters

Name	Type	Description	Mandatory	Default
host	Object name	Limit list to the ports of a specific host.	N	All hosts.
fc_host_port	N/A	Limit list to this specific port.	N	All ports
module	N/A	Limits output only to the enabled connectivity to this module.	N	All modules

Name	Type	Description	Mandatory	Default
fcport	N/A	Limits output to a specific storage system's port.	N	All ports
domain	Object name	The domain name.	N	All Domains

This command shows the connectivity status between a storage system port and a defined host. The output can be limited to a specific port, module or storage system port. Hosts can attach to the FC and iSCSI either directly (point-to-point), via a Fibre Channel fabric, or via a Gigabit Ethernet switch. Connectivity refers to both physical connectivity and SCSI login. Each output line contains the following information:

- Host (name)
- Host port (WWPN)
- Module ID, preceded by the rack ID
- Port number (within the module)

Example:

```
host_connectivity_list host=demo_host_fc0 fc_host_port=1:FC_Port:5:1
```

Output:

```
Host          Host Port        Module      Local FC port   Local iSCSI port   Type
-----         ----------       ----------- --------------- -----------------  -----
demo_host_fc0 100000062B151A98 1:Module:5  1:FC_Port:5:1                      FC
```

Field ID	Field output	Default position
host	Host	1
host_port	Host Port	2
module	Module	3
local_fc_port	Local FC port	4
local_iscsi_port	Local iSCSI port	5
type	Type	6
vlan_id	VLAN ID	7
local_ipinterface	Local IP interface	8
local_ethernet_port	Local Ethenet port	9

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Flash enclosure maintenance commands

This section describes the command-line interface (CLI) for maintaining flash enclosures.

- **Listing Flash enclosure status**
  Use the **flash_enclosure_list** command to list special flash enclosure statuses.
- **Listing a Flash canister status**
  Use the **flash_canister_list** command to list special Flash canister statuses.
- **Listing a Flash card status**
  Use the **flash_card_list** command to list special Flash card statuses.
- **Retrieving the list of Flash fans**
  Use the **flash_fan_list** command to retrieve the list of Flash fans.
- **Retrieving the list of Flash IB adapters**
  Use the **flash_ib_adapter_list** command to retrieve the list of Flash IB adapters.
- **Retrieving the Flash control connectivity list**
  Use the **flash_control_connectivity_list** command to retrieve the Flash control connectivity list.
- **Retrieving the list of Flash PSUs**
  Use the **flash_psu_list** command to retrieve the list of Flash PSUs.
- **Retrieving the list of Flash PIBs**
  Use the **flash_pib_list** command to retrieve the list of Flash PIBs.
- **Retrieving the list of Flash LED cards**
  Use the **flash_led_card_list** command to retrieve the list of Flash LED cards.
- **Listing Flash BBU components**
  Use the **flash_bbu_list** command to list Flash BBU components.
- **Retrieving the serial number of a flash enclosure 1S**
  Use the **flash_enclosure_1s_get** command to retrieve the serial number of a flash enclosure 1S.
- **Suspending the automatic Flash BBU calibration**
  Use the **flash_system_bbu_calibration_set** command to suspend and resume the automatic Flash BBU calibration.

Listing Flash enclosure status

Use the **flash_enclosure_list** command to list special flash enclosure statuses.

```
flash_enclosure_list [ flash_enclosure=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
flash_enclosure	Flash enclosure for which special statuses are to be listed.	N	All Flash Enclosures.

This command lists the status of each Flash enclosure, including:

- Component generic status
- Online canister
- Total number of canisters

Example:

```
flash_enclosure_list
```

Output:

```
Component ID           Status   Currently Functioning   Control Path Status   Cluster IP
---------------------- -------- ----------------------- --------------------- -------------
1:Flash_Enclosure:1    OK       yes                     OK                    14.10.204.3
```

```
Cont.:
```

```
Redundancy State   FW level          Has Spare   Array Rebuild Percentage
------------------ ----------------- ----------- --------------------------
online             1.5.0.0-436.13    yes         None
```

```
Cont.:
```

```
Encryption State   Machine Model
------------------ ---------------
Ready              AE2
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
control_path_status	Control Path Status	4
cluster_ip	Cluster IP	5
redundancy_state	Redundancy State	6
fw_level	FW level	7
has_spare	Has Spare	8
array_rebuild_percentage	Array Rebuild Percentage	9
machine_model	Machine Model	11
array_status	Array Status	N/A
fru_part_number	FRU Part Number	N/A
fru_identity	FRU Identity	N/A
temperature_state	Temperature State	N/A
required_service	Requires Service	N/A
service_reason	Service Reason	N/A
enabled	Enabled	N/A
cluster_id	Cluster ID	N/A
serial_number	Serial Number	N/A
encrypted	Encrypted	N/A
key_needed	Key Needed	N/A
encryption_state	Encryption State	10
base_guid	Base GUID	N/A
charging	Charging	N/A
flash_status	Flash Status	N/A
fw_upgrade_status	FW Upgrade Status	N/A
fw_upgrade_progress	FW Upgrade Progress	N/A
target_fw_version	Target FW. Version	N/A
fw_file_name	FW File Name	N/A
utility_file_name	Utility File Name	N/A
cr_key_last_modified_time	CR Key Last Modified Time	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Allowed
Read-only users	Allowed
Technicians	Allowed

Listing a Flash canister status

Use the **flash_canister_list** command to list special Flash canister statuses.

```
flash_canister_list [ canister=ComponentId | flash_enclosure=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
canister	Flash Canister for which special statuses are to be listed.	N	All Flash canisters.
flash_enclosure	Flash Enclosure for which special statuses are to be listed.	N	All Flash enclosures.

This command lists the statuses of the Flash canisters, including:

- Component generic status
- Canister ID
- Node ID
- Node name

Example:

```
flash_canister_list
```

Output:

```
Component ID            Status   Currently Functioning   Service IP      Raid Status
----------------------- -------- ----------------------- --------------- -------------
1:Flash_Canister:4:1    OK       yes                     14.10.204.77    OK
1:Flash_Canister:4:2    OK       yes                     14.10.204.110   OK
```

```
Cont.:
```

```
Control Path Status   Serial Connected   MgmtNode
--------------------- ------------------ ----------
OK                    1:Module:12        no
OK                    1:Module:13        yes
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
service_ip	Service IP	4
raid_status	Raid Status	5
control_path_status	Control Path Status	6
serial_connected	Serial Connected	7
active	MgmtNode	8
fru_part_number	FRU Part Number	N/A
fru_identity	FRU Identity	N/A
temperature_state	Temperature State	N/A
fw_level	FW Level	N/A
mac_addresses	MAC Addresses	N/A
required_service	Requires Service	N/A
service_reason	Service Reason	N/A
canister_degraded	Canister Degraded	N/A
canister_missing	Canister Missing	N/A
status_led	Status LED	N/A
check_log_led	Check Log LED	N/A
identify_led	Identify LED	N/A
controller_fault_led	Controller Fault LED	N/A
fault_led	Fault LED	N/A
dump_led	Dump LED	N/A

Field ID	Field output	Default position
canister_mode	Canister Mode	N/A
service_mode	Service Mode	N/A
miswired	Miswired	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Allowed
Read-only users	Allowed
Technicians	Allowed

Listing a Flash card status

Use the **flash_card_list** command to list special Flash card statuses.

```
flash_card_list [ flash_card=ComponentId | flash_enclosure=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
flash_card	Flash card for which special statuses are to be listed.	N	All Flash cards.
flash_enclosure	Flash enclosure for which special statuses are to be listed.	N	All Flash enclosures.

This command lists the statuses of a Flash card, including:

- Component generic status
- Slot ID
- Capacity
- Health State
- Flash type

Example:

```
flash_card_list flash_card=1:Flash_Card:4:5
```

Output:

```
Component ID       Status   Currently Functioning   Slot ID   Capacity(GB)   Health State
------------------ -------- ----------------------- --------- -------------- --------------
1:Flash_Card:4:5   OK       yes                     5         5717           good
```

```
Cont.:
```

```
Usage    Missing
-------- ---------
member   no
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
slot_id	Slot ID	4
capacity	Capacity(GB)	5
health_state	Health State	6
drive_use	Usage	7
missing	Missing	8
fru_part_number	FRU Part Number	N/A
fru_identity	FRU Identity	N/A
temperature_state	Temperature State	N/A
fw_level	FW Level	N/A
required_service	Requires Service	N/A
service_reason	Service Reason	N/A
flash_type	Type	N/A
fault_led	Fault LED	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Allowed
Read-only users	Allowed
Technicians	Allowed

Retrieving the list of Flash fans

Use the **flash_fan_list** command to retrieve the list of Flash fans.

```
flash_fan_list  [ flash_fan=ComponentId | flash_enclosure=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
flash_fan	Flash fan component ID	N	all
flash_enclosure	Flash enclosure component ID	N	all

Example:

```
flash_fan_list
```

Output:

```
Component ID       Status
------------------ --------
1:Flash_Fan:2:1    OK
1:Flash_Fan:2:2    OK
1:Flash_Fan:2:3    OK
1:Flash_Fan:2:4    OK
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
fru_part_number	FRU Part Number	N/A
fru_identity	FRU Identity	N/A
temperature_state	Temperature State	N/A
fw_level	FW Level	N/A
required_service	Requires Service	N/A
service_reason	Service Reason	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Retrieving the list of Flash IB adapters

Use the **flash_ib_adapter_list** command to retrieve the list of Flash IB adapters.

```
flash_ib_adapter_list  [ flash_ib_adapter=ComponentId | flash_enclosure=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
flash_ib_adapter	Flash IB adapter component ID	N	all
flash_enclosure	Flash Enclosure component ID	N	all

Example:

```
flash_ib_adapter_list
```

Output:

```
Component ID              Status
------------------------- --------
1:Flash_IB_Adapter:2:1    OK
1:Flash_IB_Adapter:2:2    OK
1:Flash_IB_Adapter:2:3    OK
1:Flash_IB_Adapter:2:4    OK
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
canister_id	Canister_name	4
fru_part_number	FRU Part Number	N/A
fru_identity	FRU Identity	N/A
temperature_state	Temperature State	N/A
required_service	Requires Service	N/A
service_reason	Service Reason	N/A
fw_level	FW Level	N/A
port1_id	Port 1 Component ID	N/A
port1_guid	Port 1 GUID	N/A
port2_id	Port 2 Component ID	N/A
port2_guid	Port 2 GUID	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Retrieving the Flash control connectivity list

Use the **flash_control_connectivity_list** command to retrieve the Flash control connectivity list.

```
flash_control_connectivity_list  [ canister=ComponentId | flash_enclosure=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
canister	Canister component ID	N	all
flash_enclosure	Flash enclosure Component ID	N	all

This command is used to list the Ethernet connectivity to flash canisters in the system. On FlashSystem A9000R systems, all flash canisters are connected via grid controllers (modules) 1 through 4. On FlashSystem A9000 systems, each flash canister is connected directly to a specific grid controller.

The command output lists each flash canister and its connectivity to each of the grid controllers. For each flash canister, the path to each grid controller includes the flash canister IP address used by the grid controller and the connectivity status.

Example:

```
flash_control_connectivity_list
```

Output:

```
This command's output is dependent on the FlashSystem type

Output for FlashSystem A9000R:

Component ID           Path1         PathAddr1      P1Status  Path2         PathAddr2
---------------------- ------------ ------------- ---------- ------------ --------------
1:Flash_Canister:1:1   1:Module:1   14.10.204.1    OK        1:Module:2   14.10.204.33
1:Flash_Canister:1:2   1:Module:1   14.10.204.2    OK        1:Module:2   14.10.204.34
1:Flash_Canister:2:1   1:Module:1   14.10.204.4    OK        1:Module:2   14.10.204.37
1:Flash_Canister:2:2   1:Module:1   14.10.204.5    OK        1:Module:2   14.10.204.38

P2Status   Path3         PathAddr3      P3Status   Path4         PathAddr4       P4Status
```

```
---------- ------------ -------------- ---------- ------------ --------------- ----------
OK         1:Module:3   14.10.204.65   OK         1:Module:4   14.10.204.97    OK
OK         1:Module:3   14.10.204.66   OK         1:Module:4   14.10.204.98    OK
OK         1:Module:3   14.10.204.69   OK         1:Module:4   14.10.204.101   OK
OK         1:Module:3   14.10.204.70   OK         1:Module:4   14.10.204.102   OK
```

```
Output for FlashSystem A9000:

Component ID            Path1         PathAddr1     P1Status   Path2         PathAddr2
----------------------- ------------- ------------- ---------- ------------- -------------
1:Flash_Canister:1:1                                N.A.       1:Module:2    14.10.204.1
1:Flash_Canister:1:2    1:Module:1    14.10.204.34  OK

P2Status   Path3         PathAddr3      P3Status   Path4         PathAddr4       P4Status
---------- ------------- -------------- ---------- ------------- --------------- ----------
OK                                      N.A.                                    N.A.
N.A.                                    N.A.                                    N.A.
```

Field ID	Field output	Default position
component_id	Component ID	1
virtual_ips.0.path	Path1	2
virtual_ips.0.pathAddr	PathAddr1	3
virtual_ips.0.status	P1Status	4
virtual_ips.0.state	P1State	N/A
virtual_ips.1.path	Path2	5
virtual_ips.1.pathAddr	PathAddr2	6
virtual_ips.1.status	P2Status	7
virtual_ips.1.state	P2State	N/A
virtual_ips.2.path	Path3	8
virtual_ips.2.pathAddr	PathAddr3	9
virtual_ips.2.status	P3Status	10
virtual_ips.2.state	P3State	N/A
virtual_ips.3.path	Path4	11
virtual_ips.3.pathAddr	PathAddr4	12
virtual_ips.3.status	P4Status	13
virtual_ips.3.state	P4State	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Retrieving the list of Flash PSUs

Use the **flash_psu_list** command to retrieve the list of Flash PSUs.

```
flash_psu_list  [ flash_psu=ComponentId | flash_enclosure=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
flash_psu	Flash PSU component ID	N	all
flash_enclosure	Flash Enclosure component ID	N	all

Example:

```
flash_psu_list
```

Output:

```
Component ID      Status
----------------- --------
1:Flash_PSU:2:1   OK
1:Flash_PSU:2:2   OK
```

Field ID	Field output	Default position

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
input_failed	Input Failed	4
output_failed	Output Failed	5
fru_part_number	FRU Part Number	N/A
fru_identity	FRU Identity	N/A
temperature_state	Temperature State	N/A
fw_level	FW Level	N/A
required_service	Requires Service	N/A
service_reason	Service Reason	N/A
fan_failed	Fan Failed	N/A
fault_led	Fault LED	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Retrieving the list of Flash PIBs

Use the **flash_pib_list** command to retrieve the list of Flash PIBs.

```
flash_pib_list  [ flash_pib=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
flash_pib	Flash PIB component ID	N	all

Example:

```
flash_pib_list
```

Output:

```
Component ID       Status
------------------ --------
1:Flash_PIB:2:1    OK
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
fru_part_number	FRU Part Number	N/A
fru_identity	FRU Identity	N/A
temperature_state	Temperature State	N/A
fw_level	FW Level	N/A
required_service	Requires Service	N/A
service_reason	Service Reason	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Retrieving the list of Flash LED cards

Use the **flash_led_card_list** command to retrieve the list of Flash LED cards.

```
flash_led_card_list  [ flash_led_card=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
flash_led_card	Flash LED card component ID	N	all

Example:

```
flash_led_card_list
```

Output:

```
Component ID       Status
------------------ --------
1:Flash_LED_Card:2:1   OK
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
fru_part_number	FRU Part Number	N/A
fru_identity	FRU Identity	N/A
temperature_state	Temperature State	N/A
fw_level	FW Level	N/A
required_service	Requires Service	N/A
service_reason	Service Reason	N/A
power_led	Power LED	N/A
fault_led	Fault LED	N/A
check_log_led	Check Log LED	N/A
identify_led	Identify LED	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing Flash BBU components

Use the **flash_bbu_list** command to list Flash BBU components.

```
flash_bbu_list  [ flash_bbu=ComponentId | flash_enclosure=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
flash_bbu	Flash BBU component ID	N	all
flash_enclosure	Flash Enclosure component ID	N	all

Example:

```
flash_bbu_list
```

Output:

```
Component ID       Status   Charging Status   Percent Charged   Recondition Needed
------------------ -------- ----------------- ----------------- --------------------
1:Flash_BBU:2:1    OK       idle              94                no
1:Flash_BBU:2:2    OK       idle              89                no
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
percent_charged	Percent Charged	4
recondition_needed	Recondition Needed	5
charging_status	Charging Status	6
last_recondition_time	Last Recondition Time	7
temperature_state	Temperature State	N/A
fw_level	FW Level	N/A
required_service	Requires Service	N/A
service_reason	Service Reason	N/A
fru_part_number	FRU Part Number	N/A
fru_identity	FRU Identity	N/A
fault_led	Fault LED	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Retrieving the serial number of a flash enclosure 1S

Use the **flash_enclosure_1s_get** command to retrieve the serial number of a flash enclosure 1S.

```
flash_enclosure_1s_get flash_enclosure=ComponentId
```

Parameters

Name	Description	Mandatory
flash_enclosure	Flash enclosure component ID	Y

Example:

```
flash_enclosure_1s_get
```

Field ID	Field output	Default position
full_serial	Serial	1
mtm	Mtm	N/A
serial	Serial	N/A

Example:

```
flash_enclosure_1s_get
```

Output:

```
Serial
----------------
98354151234567
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Return codes

- COMMAND_IS_NOT_VALID_DURING_FLASH_FW_UPDATE
  The requested command cannot be invoked while a FlashSystem update is running.

- FLASH_ENCLOSURE_DOES_NOT_EXIST
  Flash enclosure does not exist in the system

Suspending the automatic Flash BBU calibration

Use the **flash_system_bbu_calibration_set** command to suspend and resume the automatic Flash BBU calibration.

```
flash_system_bbu_calibration_set suspend=<yes|no> [ suspend_days=days ]
```

Parameters

Name	Type	Description	Mandatory	Default
suspend	Boolean	Suspend calibration	Y	N/A
suspend_days	Positive integer	Days to suspend	N	14

Example:

```
flash_system_bbu_calibration_set suspend=yes suspend_days=14
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Allowed

Hardware maintenance commands

This section describes the command-line interface (CLI) for maintaining hardware components.

- **Monitoring the rebuilding or redistribution processes**
  Use the **monitor_redist** command to monitor the status of the rebuilding or redistribution process.
- **Displaying the system's average power consumption**
  Use the **system_average_power_consumption** command to display the system's average power consumption.
- **Getting the values for calculating the system's average power consumption**
  Use the **system_average_power_prepare** command to fetch all the values needed to calculate the system's average power consumption.
- **Displaying the system's average temperature**
  Use the **system_average_temperature** command to display the system's average temperature.
- **Enabling XIV Support access**
  Use the **xiv_support_enable** command to enable XIV Support access for a specific period of time limiting access from the specific address.
- **Disabling XIV Support access**
  Use the **xiv_support_disable** command to disable XIV Support access.
- **Displaying the XIV Support window**
  Use the **xiv_support_show** command to display the XIV Support window.
- **Listing system components**
  Use the **component_list** command to list system components and their status.
- **Listing module configuration**
  Use the **module_list** command to list the configuration of all or specified modules.
- **Listing the internal temperature of modules**
  Use the **module_temperature_list** command to list the modules' internal temperatures in the storage system.
- **Listing boot media devices in the system**
  Use the **boot_media_list** command to lists boot media devices in the storage system.
- **Listing vault devices in the system**
  Use the **vault_device_list** command to lists special vault device statuses.
- **Listing BBUs in the system**
  Use the **module_bbu_list** command to list the BBUs in the storage system modules.
- **Listing PSUs in the system**
  Use the **module_psu_list** command to list PSUs in the system.
- **Listing compression adapters in the system**
  Use the **compression_adapter_list** command to list compression adapters in the system.
- **Listing fans in a module**
  Use the **fan_list** command to list fans in the specified module.
- **Listing NICs in the system**
  Use the **nic_list** command to list the NICs in the storage system.

- **Listing DIMMs in the modules**
  Use the **dimm_list** command to list the DIMMs in the modules.
- **Listing CPUs in the modules**
  Use the **cpu_list** command to list the central processing units (CPU) in the modules.
- **Listing InfiniBand HCA Adapters in the storage system**
  Use the **hca_list** command to list the InfiniBand host card adapters (HCAs) in the storage system.
- **Listing CNA adapters in the system**
  Use the **cna_list** command to list CNA adapters in the storage system.
- **Listing module LEDs in the system**
  Use the **module_led_list** command to display the module LED state in the system.
- **Listing data disk devices in the system**
  Use the **disk_list** command to lists special vault device statuses.
- **Listing service statuses**
  Use the **service_list** command to list all service specific statuses.
- **Listing system components that require service**
  Use the **component_service_required_list** command to list system components and their status.
- **Listing trace snapshot on a module.**
  Use the **traces_snapshot_list** command to list trace snapshots on a module.
- **Creating a trace snapshot**
  Use the **traces_snapshot** command to create a trace snapshot.

Monitoring the rebuilding or redistribution processes

Use the **monitor_redist** command to monitor the status of the rebuilding or redistribution process.

```
monitor_redist
```

This command outputs the current redistribution process running on a module. The command does not yield information about a Flash enclosure.

The command may inform you that no such process exists. If such a process exists, the following information is shown:

- Type (adding new capacity, replacing failed component, phase-out, rebuild after failure)
- Initial capacity to copy
- Time started
- Capacity remaining to copy
- Time elapsed
- Percent completed
- Estimated time to completion

Field ID	Field output	Default position
type	Type	1
partitions_total	Initial Partitions	2
partitions_left	Partitions Remaining	3
percent_done	% Done	4
time_started	Time Started	5
estimated_time_to_finish	Estimated Time to Finish	6
time_elapsed	Time Elapsed	7

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Displaying the system's average power consumption

Use the **system_average_power_consumption** command to display the system's average power consumption.

```
system_average_power_consumption
```

Field ID	Field output	Default position
value	Value	1
timestamp	Timestamp	2

Note:
The average power consumption value is indicated in Watts.

Example:

```
system_average_power_consumption
```

Output:

```
Value       Timestamp
----------- -----------
1337W       17-Jan-2020 12:30:27
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Getting the values for calculating the system's average power consumption

Use the **system_average_power_prepare** command to fetch all the values needed to calculate the system's average power consumption.

```
system_average_power_prepare
```

Example:

```
system_average_power_prepare
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Allowed

Return codes

- HOT_UPGRADE_IS_IN_PROGRESS
  The command is not allowed during hot upgrade.

  Troubleshooting: Wait for the hot upgrade to complete and try again.
- ALREADY_GETTING_FLASH_CANISTER_SNAP
  Already getting a flash canister snap.
- FAILED_GETTING_FLASH_CANISTER_SNAP
  Failed getting a flash canister snap.
- COMPONENT_DOES_NOT_EXIST
  The component does not exist.
- COMMAND_IS_NOT_VALID_DURING_FLASH_FW_UPDATE
  The requested command cannot be invoked while a FlashSystem update is running.
- SYSTEM_POWER_PREPARE_ALREADY_IN_PROGRESS
  There is already a system power prepare command in progress.

Displaying the system's average temperature

Use the **system_average_temperature** command to display the system's average temperature.

`system_average_temperature`

Field ID	Field output	Default position
value	Value	1
timestamp	Timestamp	2

Note:
The average temperature value is indicated in Celsius.

Example:

```
system_average_temperature
```

Output:

```
Value       Timestamp
----------- --------------------
17C         6-Feb-2020 12:15:22
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Enabling XIV Support access

Use the **xiv_support_enable** command to enable XIV Support access for a specific period of time limiting access from the specific address.

```
xiv_support_enable [ start=TimeStamp ] < finish=TimeStamp | timeout=Timeout > from=<IPaddress1[,IPaddress2]
...> comment=Comment
```

Parameters

Name	Type	Description	Mandatory	Default
start	N/A	Start time for allowing XIV Support access.	N	Immediately.
finish	N/A	End time for allowing XIV Support access.	N	N/A
timeout	N/A	Timeout for allowing XIV Support access in either hh:mm format, or a number of minutes. The timeout cannot exceed 23 hours and 59 minutes. The word `unlimited` denotes unexpired timeout.	N	N/A
from	N/A	The source address to which XIV Support access is limited. It may be either IPv4 or IPv6 address, or `any`, or `technician` denoting notebook port.	Y	N/A
comment	String	Reason why XIV Support access is enabled.	Y	N/A

This command enables XIV Support access for a specific period of time limiting access from the specific address.

Example:

```
xiv_support_enable finish=2012-2-3.16:30 from=192.0.2.1 comment="Some reason"
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_ENABLE_XIV_SUPPORT
  Are you sure you want to enable XIV support?

Return codes

- XIV_SUPPORT_WORK_INVALID_TIMEOUT
  Timeout must be a positive number and define a time greater then the current time
- XIV_SUPPORT_WORK_INVALID_FINISH
  The end time must be greater than the start time and the current time.
- XIV_SUPPORT_WORK_INVALID_FROM
  From must be a valid IPv4 or IPv6 address.
- LIST_WITH_MIXED_IPV6_AND_IPV4_NOT_ALLOWED
  All IP addresses in the list should be of the same type - either IPv4 or IPv6.
- LIST_WITH_ANY_OPTION_AND_SPECIFIC_IP_ADDRESSES_NOT_ALLOWED
  All IP addresses in the list should be unicast or Any. Mixing unicast and Any in the same list is not allowed.

Disabling XIV Support access

Use the **xiv_support_disable** command to disable XIV Support access.

```
xiv_support_disable
```

Example:

```
xiv_support_disable
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Displaying the XIV Support window

Use the **xiv_support_show** command to display the XIV Support window.

```
xiv_support_show
```

The following information is listed:

- From (IPv4 or IPv6 addresses, or "any address", or "technician port")
- Start (timestamp or "unlimited")
- Finish (timestamp or "unlimited")
- Comment

Example:

```
xiv_support_show
```

Output:

```
Enabled   Start                Finish               Comment
--------- -------------------- -------------------- -----------
yes       2012-03-28 12:55:21  2012-03-30 00:00:00  some work

cont:
From 0   From 1   From 2   From 3   From 4
-------- -------- -------- -------- --------
192.0.2.1
```

Field ID	Field output	Default position
enabled	Enabled	1
start	Start	2

Field ID	Field output	Default position
finish	Finish	3
comment	Comment	4
from.0	From 0	5
from.1	From 1	6
from.2	From 2	7
from.3	From 3	8
from.4	From 4	9

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Allowed
Read-only users	Allowed
Technicians	Allowed

Listing system components

Use the **component_list** command to list system components and their status.

```
component_list [ component=ComponentId ] [ filter=<ALL|FAILED|NOTOK> ]
```

Parameters

Name	Type	Description	Mandatory	Default
component	N/A	Lists only this component.	N	All components.
filter	Enumeration	Filters the list to show only failed or only non-OK components.	N	ALL

Lists system components. The list can be filtered to show only a specific component, all failed components or all components in a non-OK state.

For status and configuration of specific component types, refer to the _list commands for specific components, such as: module_list or switch_list.

Example:

```
component_list
```

Output:

```
Component ID                Status         Currently Functioning
--------------------------- -------------- ----------------------
1:BBU:12:1                  Failed         no
1:Boot_Media:12:1           OK             yes
1:Boot_Media:12:2           OK             yes
1:CNA:12:1                  OK             yes
1:CNA:13:1                  OK             yes
1:CNA:8:1                   OK             yes
1:CPU:12:1                  OK             yes
1:CPU:12:2                  OK             yes
1:DIMM:12:1                 OK             yes
1:DIMM:12:10                OK             yes
1:Data:12                   OK             yes
1:Data:13                   OK             yes
1:Data:8                    OK             yes
1:Data_Reduction:12         OK             yes
1:Data_Reduction:13         OK             yes
1:Data_Reduction:8          OK             yes
1:FC_Port:12:1              OK             yes
1:FC_Port:12:2              OK             yes
1:Fan:12:1                  OK             yes
1:Fan:12:2                  OK             yes
1:Flash_BBU:4:1             OK             yes
1:Flash_BBU:4:2             OK             yes
1:Flash_Canister:4:1        OK             yes
1:Flash_Canister:4:2        OK             yes
1:Flash_Card:4:1            Failed         no
1:Flash_Enclosure:4         OK             yes
1:Flash_Fan:4:1             OK             yes
1:Flash_Fan:4:2             OK             yes
1:Flash_Fan:4:3             OK             yes
1:Flash_Fan:4:4             OK             yes
1:Flash_IB_Adapter:4:1      OK             yes
1:Flash_IB_Adapter:4:2      OK             yes
1:Flash_IB_Adapter:4:3      OK             yes
1:Flash_IB_Adapter:4:4      OK             yes
1:Flash_LED_Card:4:1        OK             yes
```

```
1:Flash_PIB:4:1              OK              yes
1:Flash_PSU:4:1              OK              yes
1:Flash_PSU:4:2              OK              yes
1:IB_FlashSystem_Port:4:1    OK              yes
1:IB_FlashSystem_Port:4:3    OK              yes
1:IB_FlashSystem_Port:4:5    OK              yes
1:IB_FlashSystem_Port:4:7    OK              yes
1:IB_Module_Port:12:1        OK              yes
```

Output:

```
Cont:
```

```
1:IB_Module_Port:12:2        Failed          no
1:IB_Module_Port:13:1        OK              yes
1:IB_Module_Port:13:2        Failed          no
1:IB_Module_Port:8:1         OK              yes
1:IB_Module_Port:8:2         Failed          no
1:IB_Switch:1                OK              yes
1:IB_Switch:2                OK              yes
1:IB_Switch_BBU:1:1          Initializing    yes
1:IB_Switch_BBU:1:2          Initializing    yes
1:IB_Switch_BBU:2:1          Initializing    yes
1:IB_Switch_BBU:2:2          Initializing    yes
1:IB_Switch_Fan:1:1          Initializing    yes
1:IB_Switch_PSU:1:1          Initializing    yes
1:IB_Switch_PSU:1:2          Initializing    yes
1:IB_Switch_PSU:2:1          Initializing    yes
1:IB_Switch_PSU:2:2          Initializing    yes
1:IB_Switch_Port:1:12        OK              yes
1:IB_Switch_Port:1:13        OK              yes
1:IB_Switch_Port:1:20        OK              yes
1:Interface:12               OK              yes
1:Interface:13               OK              yes
1:Interface:8                OK              yes
1:Module:12                  OK              yes
1:Module:13                  OK              yes
1:Module:8                   OK              yes
1:NIC:12:1                   OK              yes
1:NIC:12:2                   OK              yes
1:PSU:12:1                   OK              yes
1:PSU:12:2                   OK              yes
1:PSU:13:1                   OK              yes
1:PSU:13:2                   OK              yes
1:PSU:8:1                    OK              yes
1:PSU:8:2                    OK              yes
1:Remote:12                  OK              yes
1:Remote:13                  OK              yes
1:Remote:8                   OK              yes
1:Vault_Device:12:1          OK              yes
1:Vault_Device:12:2          OK              yes
1:Vault_Device:13:1          OK              yes
1:Vault_Device:13:2          OK              yes
1:Vault_Device:8:1           OK              yes
1:Vault_Device:8:2           OK              yes
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
requires_service	Requires Service	N/A
service_reason	Service Reason	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing module configuration

Use the **module_list** command to list the configuration of all or specified modules.

```
module_list [ module=ModuleNumber ]
```

Parameters

Name	Description	Mandatory	Default
module	Lists the configuration of the specified module.	N	All modules

The following modules are available in FlashSystem A9000 and FlashSystem A9000R:

- BBU, also referred to or known as battery module
- Boot media, also referred to or known as hard disk drive (HDD)
- Compute module, also referred to or known as grid controller or module
- IB Switch, also referred to or known as InfiniBand switch
- Module, also referred to or known as grid controller or compute module.

This command lists the following information for each module:

- Generic component status
- Module type
- Number of boot media
- Number of vault devices
- Number of FC ports
- Number of Ethernet ports for iSCSI

Additional information is available through running **module_list -t all** :

- Serial
- Original serial
- Part number
- Original part number

Note:
The temperature values are indicated in Celsius.

Example:

```
module_list
```

Output:

```
Component ID   Status   Currently Functioning   Target Status   Type
-------------- -------- ----------------------- --------------- ------------------------
1:Module:12    OK       yes                                     g4.0_compute_enclosure
1:Module:13    OK       yes                                     g4.0_compute_enclosure
1:Module:8     OK       yes                                     g4.0_compute_enclosure
```

```
Cont.:
Boot Media Disks   Vault Devices   FC Ports   iSCSI Ports   Temperature
------------------ --------------- ---------- ------------- -------------
2                  2               4          2             22
2                  2               4          2             22
2                  2               4          2             22
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
target_status	Target Status	4
type	Type	5
boot_media_disks	Boot Media Disks	6
vault_devices	Vault Devices	7
fc_port_count	FC Ports	8
ethernet_port_count	iSCSI Ports	9
temperature	Temperature	10
enclosure_id	Enclosure ID	11
avg_power	Avg Power	N/A
serial	Serial	N/A
original_serial	Original Serial	N/A
part_number	Part Number	N/A
original_part_number	Original Part Number	N/A
sas_version	SAS	N/A
infiniband_hca_version.0	InfiniBand HCA 1	N/A
infiniband_hca_version.1	InfiniBand HCA 2	N/A
cna_version.0	CNA 1	N/A
cna_version.1	CNA 2	N/A
compression_adapter_firmware.0	Compression Adapter 1	N/A
compression_adapter_firmware.1	Compression Adapter 2	N/A
requires_service	Requires Service	N/A

Field ID	Field output	Default position
service_reason	Service Reason	N/A
memory_gb	Mem	N/A
module_11s_number	11S Number	N/A
megaraid_serial_number	MegaRAID Serial	N/A
megaraid_product_name	MegaRAID Product Name	N/A
megaraid_package_version	MegaRAID Package Version	N/A
megaraid_flash_components.0	MegaRAID Flash Component 1	N/A
megaraid_flash_components.1	MegaRAID Flash Component 2	N/A
megaraid_flash_components.2	MegaRAID Flash Component 3	N/A
megaraid_flash_components.3	MegaRAID Flash Component 4	N/A
megaraid_flash_components.4	MegaRAID Flash Component 5	N/A
megaraid_flash_components.5	MegaRAID Flash Component 6	N/A
megaraid_flash_components.6	MegaRAID Flash Component 7	N/A
megaraid_flash_components.7	MegaRAID Flash Component 8	N/A
imm_version	IMM Version	N/A
uefi_version	UEFI Version	N/A
dsa_version	DSA Version	N/A
me_version	ME Version	N/A
mcu_version	MCU Version	N/A
board_serial	Board Serial	N/A
board_part_number	Board Part Number	N/A
board_mfg_date	Board MFG Date	N/A
dasd_board_serial	Backplane Serial	N/A
dasd_board_part_number	Backplane Part Number	N/A
dasd_board_manufacturer	Backplane Manufacturer	N/A
dasd_board_mfg_date	Backplane MFG Date	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing the internal temperature of modules

Use the **module_temperature_list** command to list the modules' internal temperatures in the storage system.

```
module_temperature_list [ module=ModuleNumber ]
```

Parameters

Name	Description	Mandatory	Default
module	Limits the listing to a specific module.	N	All temperatures in all modules.

Note:
The temperature values are indicated in Celsius.

Example:

```
module_temperature_list -f all
```

Output:

```
Module         Ambient   PCH   RS1   RS2   MID1   MID2   RAID   DIMM AB VR   DIMM CD VR
-------------- --------- ----- ----- ----- ------ ------ ------ ------------ ------------
1:Module:1     16        44    32    38    31     31     34     27           27
1:Module:11    16        46    35    40    34     33     35     26           28
1:Module:2     16        43    34    38    33     32     33     26           27
1:Module:4     16        43    34    37    33     32     34     25           26

cont:
DIMM EF VR   DIMM GH VR   CPU1   CPU2   InfiniBand HCA   PSU FR   Fibre Channel Port fc-0
------------ ------------ ------ ------ ---------------- -------- -------------------------
28           27           52     55     53               28       0
27           32           58     68     56               32       0
```

```
28          27          51     53     53              28       0
29          30          51     55     53              29       0
```

```
cont:
Fibre Channel Port fc-2    CNA=0   CNA-1   Boot Device 0   Boot Device 1   Vault Device 0
-------------------------  ------- ------- --------------- --------------- ----------------
0                          50      46      23              21              20
0                          51      52      27              25              21
0                          50      49      22              22              20
0                          48      52      24              23              20
```

```
cont:
Vault Device 1   BBU1   BBU2   CPU1 VR   CPU2 VR   Fibre Channel Port fc-1   Fibre Channel
                                                                            Port fc-3
---------------- ------ ------ --------- --------- -------------------------- -------------
21               18     18     35        36        0                          0
21               17     16     37        40        0                          0
20               18     17     35        37        0                          0
20               17     17     35        37        0                          0
```

Field ID	Field output	Default position
component_id	Module	1
sdr_temperatures.0	Ambient	2
sdr_temperatures.1	PCH	3
sdr_temperatures.2	RS1	4
sdr_temperatures.3	RS2	5
sdr_temperatures.4	MID1	6
sdr_temperatures.5	MID2	7
sdr_temperatures.6	RAID	8
sdr_temperatures.8	CPU1 VR	N/A
sdr_temperatures.9	CPU2 VR	N/A
sdr_temperatures.10	DIMM AB VR	9
sdr_temperatures.11	DIMM CD VR	10
sdr_temperatures.12	DIMM EF VR	11
sdr_temperatures.13	DIMM GH VR	12
sdr_temperatures.14	CPU1	13
sdr_temperatures.15	CPU2	14
ib_hca_temperature	InfiniBand HCA	15
sdr_temperatures.7	PSU FR	16
fc_adapter_temperature.0	Fibre Channel Port fc-0	17
fc_adapter_temperature.1	Fibre Channel Port fc-1	N/A
fc_adapter_temperature.2	Fibre Channel Port fc-2	18
fc_adapter_temperature.3	Fibre Channel Port fc-3	N/A
cna_temperature.0	CNA=0	19
cna_temperature.1	CNA-1	20
boot_device_temperature.0	Boot Device 0	21
boot_device_temperature.1	Boot Device 1	22
vault_device_temperature.0	Vault Device 0	23
vault_device_temperature.1	Vault Device 1	24
bbu_temperature.0	BBU1	25
bbu_temperature.1	BBU2	26

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing boot media devices in the system

Use the **boot_media_list** command to lists boot media devices in the storage system.

```
boot_media_list [ module=ModuleNumber | boot_media=BootMediaDevice ]
```

Parameters

Name	Description	Mandatory	Default
module	Limits the listing to the specific module.	N	Boot media devices in a specific module.
boot_media	Limits the listing to a specific boot media.	N	A specific boot media device.

Note:
The temperature values are indicated in Celsius.

Example:

```
boot_media_list -f all
```

Output:

```
Component ID         Status   Currently Functioning   Hardware Status   Vendor
-------------------- -------- ----------------------- ----------------- --------
1:Boot_Media:11:1    OK       yes                     OK                IBM-ESXS
1:Boot_Media:11:2    OK       yes                     OK                IBM-ESXS
1:Boot_Media:1:1     OK       yes                     OK                IBM-ESXS
1:Boot_Media:1:2     OK       yes                     OK                IBM-ESXS
1:Boot_Media:2:1     OK       yes                     OK                IBM-ESXS
1:Boot_Media:2:2     OK       yes                     OK                IBM-ESXS
1:Boot_Media:4:1     OK       yes                     OK                IBM-ESXS
1:Boot_Media:4:2     OK       yes                     OK                IBM-ESXS

Model               Serial   FW     Temperature   Original Serial   Part #
------------------- -------- ------ ------------- ----------------- --------
HUC101860CS420      03V0E75K J5H2   26            03V0E75K
HUC101860CS420      03V0LPEK J5H2   25            03V0LPEK
HUC101860CS420      03V0X90H J5H2   21            03V0X90H
HUC101860CS420      03V0AHME J5H2   20            03V0AHME
HUC101860CS420      03V117XE J5H2   22            03V117XE
HUC101860CS420      03V0DAGK J5H2   21            03V0DAGK
HUC101860CS420      03V0DJ4K J5H2   23            03V0DJ4K
HUC101860CS420      03V0H38K J5H2   22            03V0H38K

Original Part #  Size (GB)   Requires Service   Service Reason
---------------- ---------- ------------------ ----------------
                 600 GB
                 600 GB
                 600 GB
                 600 GB
                 600 GB
                 600 GB
                 600 GB
                 600 GB

Rebuild Progress   Rebuild Time (sec.)
------------------ --------------------
N/A                N/A
N/A                N/A
N/A                N/A
N/A                N/A
N/A                N/A
N/A                N/A
N/A                N/A
N/A                N/A
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
hardware_status	Hardware Status	4
vendor	Vendor	5
model	Model	6
serial	Serial	7
fw_revision	FW	8
temperature	Temperature	N/A
original_serial	Original Serial	N/A
part_number	Part #	N/A
original_part_number	Original Part #	N/A
fru_pn	FRU PN	N/A
original_fru_pn	Original FRU PN	N/A
size	Size	N/A
requires_service	Requires Service	N/A
service_reason	Service Reason	N/A
rebuild_progress	Rebuild Progress	N/A
rebuild_time	Rebuild Time	N/A
media_error_count	Media Errors	N/A

Field ID	Field output	Default position
other_error_count	Other Errors	N/A
predictive_failure_count	Predictive Failures	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing vault devices in the system

Use the **vault_device_list** command to lists special vault device statuses.

```
vault_device_list [ module=ModuleNumber  | vault_device=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
module	Limits the listing to a specific module.	N	All vault devices in all modules.
vault_device	Vault devices for which special statuses are to be listed.	N	All vault devices.

Note:
The temperature values are indicated in Celsius.

Example:

```
vault_device_list
```

Output:

```
        Component ID           Status   Currently Functioning   Capacity   Target Status   Vendor
        ---------------------- -------- ----------------------- ---------- --------------- --------
        1:Vault_Device:14:1    OK       yes                     250GB                      LENOVO-X
        1:Vault_Device:14:2    OK       yes                     250GB                      LENOVO-X
        1:Vault_Device:3:1     OK       yes                     250GB                      LENOVO-X
        1:Vault_Device:3:2     OK       yes                     250GB                      LENOVO-X
        1:Vault_Device:5:1     OK       yes                     250GB                      LENOVO-X
        1:Vault_Device:5:2     OK       yes                     250GB                      LENOVO-X
        1:Vault_Device:6:1     OK       yes                     250GB                      LENOVO-X
        1:Vault_Device:6:2     OK       yes                     250GB                      LENOVO-X
```

```
        Cont.:
        Model              Serial     Firmware   FRU       Temperature   Encryption State
        ------------------ ---------- ---------- --------- ------------- ------------------
        HUSMR1625ASS20E    0PVGJTPA   P4C9       00NA685   22            Ready
        HUSMR1625ASS20E    0PVGHN6A   P4C9       00NA685   22            Ready
        HUSMR1625ASS20E    0PVJ2PEA   P4C9       00NA685   21            Ready
        HUSMR1625ASS20E    0PVJ251A   P4C9       00NA685   21            Ready
        HUSMR1625ASS20E    0PVJ99UA   P4C9       00NA685   20            Ready
        HUSMR1625ASS20E    0PVJRB5A   P4C9       00NA685   21            Ready
        HUSMR1625ASS20E    0PVJS39A   P4C9       00NA685   20            Ready
        HUSMR1625ASS20E    0PVJ9RAA   P4C9       00NA685   21            Ready
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
capacity_in_bytes	Capacity	N/A
capacity	Capacity	4
target_status	Target Status	5
vendor	Vendor	6
original_vendor	Original Vendor	N/A
model	Model	7
original_model	Original Model	N/A
serial	Serial	8
original_serial	Original Serial	N/A
firmware	Firmware	9
original_firmware	Original Firmware	N/A

Field ID	Field output	Default position
part_number	FRU	10
original_part_number	Original FRU	N/A
temperature	Temperature	11
encryption_state	Encryption State	12
hw_mon_node_id	Hw Node Owner	N/A
requires_service	Requires Service	N/A
service_reason	Service Reason	N/A
revision	Revision	N/A
drive_pn	Drive P/N	N/A
original_drive_pn	Original Drive P/N	N/A
fru_pn	FRU P/N	N/A
original_fru_pn	Original FRU P/N	N/A
desc.bgd_scan	Background Scan	N/A
desc.disk_id	Device ID	N/A
desc.last_sample_serial	Last Sample Serial	N/A
desc.last_sample_time	Last Sample Time	N/A
desc.power_is_on	Power On	N/A
desc.power_on_hours	Power On Hours	N/A
desc.power_on_minutes	Power On Minutes	N/A
desc.last_time_pom_was_mod	Last Time Power On Minutes Was Modified	N/A
desc.read_fail	Read Fail	N/A
desc.smart_code	SMART Code	N/A
desc.smart_fail	SMART Fail	N/A
desc.temperature_status.reported_severity	Reported Temperature Severity	N/A
desc.temperature_status.reported_temperature	Reported Temperature	N/A
desc.temperature_status.temperature	Device Temperature	N/A
desc.sw_encryption_active	Software-Based Encryption Active	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Allowed
Read-only users	Allowed
Technicians	Allowed

Listing BBUs in the system

Use the **module_bbu_list** command to list the BBUs in the storage system modules.

```
module_bbu_list [ module=ModuleNumber | module_bbu=BbuNumber ]
```

Parameters

Name	Description	Mandatory	Default
module	Limits the listing to a specific module.	N	All BBUs in all modules.
module_bbu	Limits the listing to a specific bbu.	N	ALL BBUs.

Example:

```
module_bbu_list -f all
```

Output:

```
Component ID   Status   Curr Functioning   State   Hardware Status   Remaining Capacity
-------------- -------- ------------------ ------- ----------------- --------------------
1:BBU:14:1     OK       yes                Full    OK                891
1:BBU:14:2     OK       yes                Full    OK                877
1:BBU:3:1      OK       yes                Full    OK                787
1:BBU:3:2      OK       yes                Full    OK                860
1:BBU:5:1      OK       yes                Full    OK                792
1:BBU:5:2      OK       yes                Full    OK                898
1:BBU:6:1      OK       yes                Full    OK                817
1:BBU:6:2      OK       yes                Full    OK                814
```

```
cont:
Full Charge Capacity   Charged %   Time to Empty   Time to Full   Charger State
```

```
---------------------- ----------- --------------- -------------- ---------------
891                    100         1600200         0              in progress
877                    100         0               0              in progress
787                    100         0               0              in progress
860                    100         0               0              in progress
792                    100         0               0              in progress
898                    100         0               0              in progress
817                    100         2944800         0              in progress
814                    100         0               0              in progress
```

```
cont:
Calibration State   Calibration Time
------------------- ------------------
Idle                0
Idle                0
Idle                0
Idle                0
Idle                0
Idle                0
Idle                0
Idle                0
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Curr Functioning	3
required_service	Requires Service	N/A
service_reason	Service Reason	N/A
bbu_state	State	4
bbu_status	Hardware Status	5
remaining_capacity	Remaining Capacity	6
full_charge_capacity	Full Charge Capacity	7
percent_charged	Charged %	8
time_to_empty	Time to Empty	9
time_to_full	Time to Full	10
charger_state	Charger State	11
calib_state	Calibration State	12
calib_requested	Calibration Needed	N/A
last_succ_calib_date	Successful Calibration Time	N/A
last_calib_date	Calibration Time	13
last_calib_result	Calibration Result	N/A
insertion_date	Inserted	N/A
manuf_date	Manufactured	N/A
fw	FW Version	N/A
epow_cable_present	EPOW Cable Present	N/A
power_sense_cable_present	Power Sense Cable Present	N/A
epow_simulate	EPOW Simulate	N/A
epow_asserted	EPOW Asserted	N/A
cycle_count	Cycles	N/A
temperature_tenths_celsius	Temp /10C	N/A
charger_enabled	Charger Enabled	N/A
slow_charge_enabled	Slow Charge Enabled	N/A
discharge_enabled	Discharge Enabled	N/A
ps2_present	PS2 Present	N/A
charge_now	Nominal Available Capacity mAh	N/A
voltage_now	Voltage Now mV	N/A
current_now	Current Now mA	N/A
power_avg	Power Average mW	N/A
charge_full	Full Available Capacity mAh	N/A
charge_full_design	Design Charge	N/A
energy_now	Energy now mWh	N/A
at_rate	At Rate	N/A
at_rate_tte	At Rate Time to Empty	N/A
charge_now_sufficient	Charge Now Sufficient	N/A
endurance_start_monotonic_time	Endurance Start Monotonic Time	N/A
serial	Serial	N/A
original_serial	Original Serial	N/A
part_number	Part #	N/A
original_part_number	Original Part #	N/A
fru	FRU	N/A

Field ID	Field output	Default position
runtime	Runtime	N/A
full_power_runtime	Full Power Runtime	N/A
half_power_runtime	Half Power Runtime	N/A
module_runtime	Module Runtime	N/A
state_of_health	Health	N/A
charge_voltage	Charge Voltage mV	N/A
charge_current	Charge Current mA	N/A
test_calib_en	Test/Calib. Enabled	N/A
fhd_enabled	FHD Enabled	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing PSUs in the system

Use the **module_psu_list** command to list PSUs in the system.

```
module_psu_list [ module=ModuleNumber  | psu=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
module	Limits the listing to a specific module.	N	All PSUs in all modules.
psu	Lists only a specific PSU.	N	A specific PSU.

Example:

```
module_psu_list -f all
```

Output:

```
Component ID   Status   Currently Functioning   Location
-------------- -------- ----------------------- ----------------
1:PSU:1:1      Failed   no                      Power Supply 1
1:PSU:1:2      OK       yes                     Power Supply 2
1:PSU:2:1      OK       yes                     Power Supply 1
1:PSU:2:2      Failed   no                      Power Supply 2
1:PSU:4:1      OK       yes                     Power Supply 1
1:PSU:4:2      Failed   no                      Power Supply 2

cont:
Sensor status                              Serial number   Part number
------------------------------------------ --------------- -------------
Presence NOT detected                      N/A             N/A
Presence detected                          K115148J01L     94Y8143
Presence detected                          K115148B06E     94Y8143
Presence detected, Power Supply AC lost    K115148B01P     94Y8143
Presence detected                          K115148B0AP     94Y8143
Presence detected, Power Supply AC lost    K115148J01B     94Y8143

cont:
Requires Service   Service Reason
------------------ -------------------------
REPLACE            MODULE_PSU__NOT_DETECTED


COMPONENT_TEST     MODULE_PSU__BAD_POWER_INPUT

COMPONENT_TEST     MODULE_PSU__BAD_POWER_INPUT
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
location	Location	4
sensor_statuses	Sensor statuses	5

Field ID	Field output	Default position
requires_service	Requires Service	N/A
service_reason	Service Reason	N/A
serial	Serial number	N/A
last_valid_serial	Last Valid Serial number	N/A
part_number	Part number	N/A
manufacturer	Manufacturer	N/A
mfg_date	Manufacturing Date	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing compression adapters in the system

Use the **compression_adapter_list** command to list compression adapters in the system.

```
compression_adapter_list [ module=ModuleNumber  | compression_adapter=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
module	Limits the listing to a specific module.	N	All compression adapters in all modules.
compression_adapter	Lists only a specific compression adapter.	N	A specific compression adapter.

Note: The Technician Assistant tool and XCLI Utility refer to the Ethernet adapter as compression adapter.
Compression adapters are used to increase the speed of I/O transfers to and from compressed volumes.
Example:

```
compression_adapter_list
```

Output:

```
Component ID                 Status   Currently Functioning   Firmware   Hardware
--------------------------- -------- ----------------------- ---------- ----------
1:Compression_Adapter:6:1    OK       yes                     1.0.12     A0 SKU3
1:Compression_Adapter:6:2    OK       yes                     1.0.12     A0 SKU3

Driver
----------------
1.0.12 3d60d12
1.0.12 3d60d12
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
serial	Serial	4
firmware_version	Firmware	5
hardware_version	Hardware	6
driver_version	Driver	7
type	Type	N/A
mmp_version	MMP	N/A
pci_address	PCI Address	N/A
quick_assist_api_cy_version	QuickAssist API CY	N/A
quick_assist_api_dc_version	QuickAssist API DC	N/A
threading_mode	Threading Mode	N/A
requires_service	Requires Service	N/A
service_reason	Service Reason	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing fans in a module

Use the **fan_list** command to list fans in the specified module.

```
fan_list [ module=ModuleNumber  | fan=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
module	Limits the listing to a specific module.	N	All fans in all modules.
fan	Lists only a specific fan.	N	A specific fan.

Example:

```
fan_list
```

Output:

```
Component ID   Status   Currently Functioning   Location   Speed
-------------- -------- ----------------------- ---------- -------
1:Fan:14:1     OK       yes                     1A         4012
1:Fan:14:10    OK       yes                     5B         3780
1:Fan:14:11    OK       yes                     6A         3953
1:Fan:14:12    OK       yes                     6B         3780
1:Fan:14:2     OK       yes                     1B         3717
1:Fan:14:3     OK       yes                     2A         4012
1:Fan:14:4     OK       yes                     2B         3780
1:Fan:14:5     OK       yes                     3A         4012
1:Fan:14:6     OK       yes                     3B         3780
1:Fan:14:7     OK       yes                     4A         4012
1:Fan:14:8     OK       yes                     4B         3780
1:Fan:14:9     OK       yes                     5A         3894
1:Fan:3:1      OK       yes                     1A         3894
1:Fan:3:10     OK       yes                     5B         3717
1:Fan:3:11     OK       yes                     6A         3953
1:Fan:3:12     OK       yes                     6B         3843
1:Fan:3:2      OK       yes                     1B         3717
1:Fan:3:3      OK       yes                     2A         4012
1:Fan:3:4      OK       yes                     2B         3780
1:Fan:3:5      OK       yes                     3A         3953
1:Fan:3:6      OK       yes                     3B         3780
1:Fan:3:7      OK       yes                     4A         3953
1:Fan:3:8      OK       yes                     4B         3717
1:Fan:3:9      OK       yes                     5A         3894
1:Fan:5:1      OK       yes                     1A         3953
1:Fan:5:10     OK       yes                     5B         3528
1:Fan:5:11     OK       yes                     6A         3953
1:Fan:5:12     OK       yes                     6B         3780
1:Fan:5:2      OK       yes                     1B         3780
1:Fan:5:3      OK       yes                     2A         3953
1:Fan:5:4      OK       yes                     2B         3780
1:Fan:5:5      OK       yes                     3A         3894
1:Fan:5:6      OK       yes                     3B         3780
1:Fan:5:7      OK       yes                     4A         3953
1:Fan:5:8      OK       yes                     4B         3780
1:Fan:5:9      OK       yes                     5A         3953
1:Fan:6:1      OK       yes                     1A         3953
1:Fan:6:10     OK       yes                     5B         3780
1:Fan:6:11     OK       yes                     6A         3894
1:Fan:6:12     OK       yes                     6B         3717
1:Fan:6:2      OK       yes                     1B         3780
1:Fan:6:3      OK       yes                     2A         3953
1:Fan:6:4      OK       yes                     2B         3654
1:Fan:6:5      OK       yes                     3A         3953
1:Fan:6:6      OK       yes                     3B         3843
1:Fan:6:7      OK       yes                     4A         3953
1:Fan:6:8      OK       yes                     4B         3780
1:Fan:6:9      OK       yes                     5A         4012
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2

Field ID	Field output	Default position
currently_functioning	Currently Functioning	3
location_a	Location	N/A
rpm_a	Speed	4
min_rpm_a	Min Speed	N/A
max_rpm_a	Max Speed	N/A
location_b	Peer Location	N/A
rpm_b	Peer Speed	5
rpm_b	Peer Speed	N/A
min_rpm_b	Peer Min Speed	N/A
max_rpm_b	Peer Max Speed	N/A
requires_service	Requires Service	N/A
service_reason	Service Reason	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing NICs in the system

Use the **nic_list** command to list the NICs in the storage system.

```
nic_list [ module=ModuleNumber  | nic=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
module	Limits the listing to a specific module.	N	All NICs in all modules.
nic	Lists only a specific NIC.	N	A specific NIC.

Example:

```
nic_list -f all
```

Output:

```
Component ID   Status   Currently Functioning   Hardware Status   Device Name
-------------- -------- ----------------------- ----------------- -------------
1:NIC:10:1     OK       yes                     OK                eth0
1:NIC:10:2     OK       yes                     OK                eth1
1:NIC:10:3     OK       yes                     OK                eth2
1:NIC:10:4     OK       yes                     OK                eth3
1:NIC:10:5     OK       yes                     OK                eth4
1:NIC:10:6     OK       yes                     OK                eth5
1:NIC:11:1     OK       yes                     OK                eth0
1:NIC:11:2     OK       yes                     OK                eth1
1:NIC:11:3     OK       yes                     OK                eth2
1:NIC:11:4     OK       yes                     OK                eth3
1:NIC:11:5     OK       yes                     OK                eth4
1:NIC:11:6     OK       yes                     OK                eth5
1:NIC:7:1      OK       yes                     OK                eth0
1:NIC:7:2      OK       yes                     OK                eth1
1:NIC:7:3      OK       yes                     OK                eth2
1:NIC:7:4      OK       yes                     OK                eth3
1:NIC:9:1      OK       yes                     OK                eth0
1:NIC:9:2      OK       yes                     OK                eth1
1:NIC:9:3      OK       yes                     OK                eth2
1:NIC:9:4      OK       yes                     OK                eth3
1:NIC:9:5      OK       yes                     OK                eth4
1:NIC:9:6      OK       yes                     OK                eth5
```

```
Cont.:
```

```
Serial              Original Serial     Part #
------------------- ------------------- -------------------------------------
40:f2:e9:af:26:b0   40:f2:e9:af:26:b0   14e4_1657_40f2e9af26b0_5719-v1.38
40:f2:e9:af:26:b1   40:f2:e9:af:26:b1   14e4_1657_40f2e9af26b1_5719-v1.38
40:f2:e9:af:26:b2   40:f2:e9:af:26:b2   14e4_1657_40f2e9af26b2_5719-v1.38
40:f2:e9:af:26:b3   40:f2:e9:af:26:b3   14e4_1657_40f2e9af26b3_5719-v1.38
```

```
f4:52:14:6e:8f:70   f4:52:14:6e:8f:70   15b3_1007_708f6e00031452f4_2.35.5100
f4:52:14:6e:8f:71   f4:52:14:6e:8f:71   15b3_1007_708f6e00031452f4_2.35.5100
40:f2:e9:af:24:48   40:f2:e9:af:24:48   14e4_1657_40f2e9af2448_5719-v1.38
40:f2:e9:af:24:49   40:f2:e9:af:24:49   14e4_1657_40f2e9af2449_5719-v1.38
40:f2:e9:af:24:4a   40:f2:e9:af:24:4a   14e4_1657_40f2e9af244a_5719-v1.38
40:f2:e9:af:24:4b   40:f2:e9:af:24:4b   14e4_1657_40f2e9af244b_5719-v1.38
f4:52:14:6e:8d:30   f4:52:14:6e:8d:30   15b3_1007_308d6e00031452f4_2.35.5100
f4:52:14:6e:8d:31   f4:52:14:6e:8d:31   15b3_1007_308d6e00031452f4_2.35.5100
40:f2:e9:af:2a:90   40:f2:e9:af:2a:90   14e4_1657_40f2e9af2a90_5719-v1.38
40:f2:e9:af:2a:91   40:f2:e9:af:2a:91   14e4_1657_40f2e9af2a91_5719-v1.38
40:f2:e9:af:2a:92   40:f2:e9:af:2a:92   14e4_1657_40f2e9af2a92_5719-v1.38
40:f2:e9:af:2a:93   40:f2:e9:af:2a:93   14e4_1657_40f2e9af2a93_5719-v1.38
40:f2:e9:af:23:b8   40:f2:e9:af:23:b8   14e4_1657_40f2e9af23b8_5719-v1.38
40:f2:e9:af:23:b9   40:f2:e9:af:23:b9   14e4_1657_40f2e9af23b9_5719-v1.38
40:f2:e9:af:23:ba   40:f2:e9:af:23:ba   14e4_1657_40f2e9af23ba_5719-v1.38
40:f2:e9:af:23:bb   40:f2:e9:af:23:bb   14e4_1657_40f2e9af23bb_5719-v1.38
f4:52:14:6e:8f:20   f4:52:14:6e:8f:20   15b3_1007_208f6e00031452f4_2.35.5100
f4:52:14:6e:8f:21   f4:52:14:6e:8f:21   15b3_1007_208f6e00031452f4_2.35.5100
```

```
Cont.:

Original Part Number                    Requires Service   Service Reason
--------------------------------------- ------------------ ----------------
14e4_1657_40f2e9af26b0_5719-v1.38
14e4_1657_40f2e9af26b1_5719-v1.38
14e4_1657_40f2e9af26b2_5719-v1.38
14e4_1657_40f2e9af26b3_5719-v1.38
15b3_1007_708f6e00031452f4_2.35.5100
15b3_1007_708f6e00031452f4_2.35.5100
14e4_1657_40f2e9af2448_5719-v1.38
14e4_1657_40f2e9af2449_5719-v1.38
14e4_1657_40f2e9af244a_5719-v1.38
14e4_1657_40f2e9af244b_5719-v1.38
15b3_1007_308d6e00031452f4_2.35.5100
15b3_1007_308d6e00031452f4_2.35.5100
14e4_1657_40f2e9af2a90_5719-v1.38
14e4_1657_40f2e9af2a91_5719-v1.38
14e4_1657_40f2e9af2a92_5719-v1.38
14e4_1657_40f2e9af2a93_5719-v1.38
14e4_1657_40f2e9af23b8_5719-v1.38
14e4_1657_40f2e9af23b9_5719-v1.38
14e4_1657_40f2e9af23ba_5719-v1.38
14e4_1657_40f2e9af23bb_5719-v1.38
15b3_1007_208f6e00031452f4_2.35.5100
15b3_1007_208f6e00031452f4_2.35.5100
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
hardware_status	Hardware Status	4
device_name	Device Name	5
serial	Serial	N/A
original_serial	Original Serial	N/A
part_number	Part #	N/A
original_part_number	Original Part Number	N/A
requires_service	Requires Service	N/A
service_reason	Service Reason	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing DIMMs in the modules

Use the **dimm_list** command to list the DIMMs in the modules.

```
dimm_list [ module=ModuleNumber  | dimm=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
module	Limits the listing to a specific module.	N	All DIMMs in all modules.
dimm	Lists only a specific DIMM.	N	A specific DIMM.

The memory modules (DIMMs) run the microcode and the data cache in the grid controller.

Example:

```
dimm_list -f all
```

Output:

```
Component ID   Status   Currently Functioning   Hardware Status   DIMM Id   CPU
-------------- -------- ----------------------- ----------------- --------- -----
1:DIMM:7:1     OK       yes                     OK                1         1
1:DIMM:7:10    OK       yes                     OK                10        1
1:DIMM:7:11    OK       yes                     OK                11        1
1:DIMM:7:12    OK       yes                     OK                12        1
1:DIMM:7:13    OK       yes                     OK                13        2
1:DIMM:7:14    OK       yes                     OK                14        2
1:DIMM:7:15    OK       yes                     OK                15        2
1:DIMM:7:16    OK       yes                     OK                16        2
1:DIMM:7:17    OK       yes                     OK                17        2
1:DIMM:7:18    OK       yes                     OK                18        2
1:DIMM:7:19    OK       yes                     OK                19        2
1:DIMM:7:2     OK       yes                     OK                2         1
1:DIMM:7:20    OK       yes                     OK                20        2
1:DIMM:7:21    OK       yes                     OK                21        2
1:DIMM:7:22    OK       yes                     OK                22        2
1:DIMM:7:23    OK       yes                     OK                23        2
1:DIMM:7:24    OK       yes                     OK                24        2
1:DIMM:7:3     OK       yes                     OK                3         1
1:DIMM:7:4     OK       yes                     OK                4         1
1:DIMM:7:5     OK       yes                     OK                5         1
1:DIMM:7:6     OK       yes                     OK                6         1
1:DIMM:7:7     OK       yes                     OK                7         1
1:DIMM:7:8     OK       yes                     OK                8         1
1:DIMM:7:9     OK       yes                     OK                9         1
1:DIMM:9:1     OK       yes                     OK                1         1
1:DIMM:9:10    OK       yes                     OK                10        1
1:DIMM:9:11    OK       yes                     OK                11        1
1:DIMM:9:12    OK       yes                     OK                12        1
1:DIMM:9:13    OK       yes                     OK                13        2
1:DIMM:9:14    OK       yes                     OK                14        2
1:DIMM:9:15    OK       yes                     OK                15        2
1:DIMM:9:16    OK       yes                     OK                16        2
1:DIMM:9:17    OK       yes                     OK                17        2
1:DIMM:9:18    OK       yes                     OK                18        2
1:DIMM:9:19    OK       yes                     OK                19        2
1:DIMM:9:2     OK       yes                     OK                2         1
1:DIMM:9:20    OK       yes                     OK                20        2
1:DIMM:9:21    OK       yes                     OK                21        2
1:DIMM:9:22    OK       yes                     OK                22        2
1:DIMM:9:23    OK       yes                     OK                23        2
1:DIMM:9:24    OK       yes                     OK                24        2
1:DIMM:9:3     OK       yes                     OK                3         1
1:DIMM:9:4     OK       yes                     OK                4         1
1:DIMM:9:5     OK       yes                     OK                5         1
1:DIMM:9:6     OK       yes                     OK                6         1
1:DIMM:9:7     OK       yes                     OK                7         1
1:DIMM:9:8     OK       yes                     OK                8         1
1:DIMM:9:9     OK       yes                     OK                9         1
```

Cont.:

Output:

```
Size(Mb)   Speed(MHz)   Configured Clock Speed(MHz)   Manufacturer   Serial
---------- ------------ ----------------------------- -------------- ----------
16384      2133         2133                          Hynix          505F63F9
16384      2133         2133                          Samsung        39542977
16384      2133         2133                          Samsung        3953EA6C
16384      2133         2133                          Samsung        3953F39A
16384      2133         2133                          Samsung        3953F240
16384      2133         2133                          Samsung        39542562
16384      2133         2133                          Samsung        3954300B
16384      2133         2133                          Samsung        39546472
16384      2133         2133                          Samsung        39540BB2
16384      2133         2133                          Samsung        3953FB59
16384      2133         2133                          Samsung        3954074A
16384      2133         2133                          Samsung        3953F241
16384      2133         2133                          Samsung        395404E0
16384      2133         2133                          Samsung        395425D8
16384      2133         2133                          Samsung        39542BF2
16384      2133         2133                          Samsung        395426EF
16384      2133         2133                          Samsung        3953EB61
16384      2133         2133                          Samsung        39542AD0
16384      2133         2133                          Samsung        39542973
```

```
16384      2133         2133                        Samsung        39542ACF
16384      2133         2133                        Samsung        3953E982
16384      2133         2133                        Samsung        39542568
16384      2133         2133                        Samsung        3953EA4A
16384      2133         2133                        Samsung        3953E993
16384      2133         2133                        Hynix          804AC8C2
16384      2133         2133                        Hynix          707387FA
16384      2133         2133                        Hynix          7073895C
16384      2133         2133                        Hynix          3077315A
16384      2133         2133                        Hynix          90655EDF
16384      2133         2133                        Hynix          7073885E
16384      2133         2133                        Hynix          30772789
16384      2133         2133                        Hynix          707388BB
16384      2133         2133                        Hynix          90655FC6
16384      2133         2133                        Hynix          70738960
16384      2133         2133                        Hynix          70738871
16384      2133         2133                        Hynix          7073881A
16384      2133         2133                        Hynix          307C97D6
16384      2133         2133                        Hynix          7073880D
16384      2133         2133                        Hynix          70738819
16384      2133         2133                        Hynix          30773136
16384      2133         2133                        Hynix          30772FF8
16384      2133         2133                        Hynix          70738850
16384      2133         2133                        Hynix          3079B2CD
16384      2133         2133                        Hynix          307A04D2
16384      2133         2133                        Hynix          70738951
16384      2133         2133                        Hynix          70738939
16384      2133         2133                        Hynix          70738A0B
16384      2133         2133                        Hynix          70738823
```

Cont.:

Output:

```
Original          Serial Part #       Original Part Number
----------------- ------------------  ----------------------
505F63F9          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
39542977          M393A2G40DB0-CPB    M393A2G40DB0-CPB
3953EA6C          M393A2G40DB0-CPB    M393A2G40DB0-CPB
3953F39A          M393A2G40DB0-CPB    M393A2G40DB0-CPB
3953F240          M393A2G40DB0-CPB    M393A2G40DB0-CPB
39542562          M393A2G40DB0-CPB    M393A2G40DB0-CPB
3954300B          M393A2G40DB0-CPB    M393A2G40DB0-CPB
39546472          M393A2G40DB0-CPB    M393A2G40DB0-CPB
39540BB2          M393A2G40DB0-CPB    M393A2G40DB0-CPB
3953FB59          M393A2G40DB0-CPB    M393A2G40DB0-CPB
3954074A          M393A2G40DB0-CPB    M393A2G40DB0-CPB
3953F241          M393A2G40DB0-CPB    M393A2G40DB0-CPB
395404E0          M393A2G40DB0-CPB    M393A2G40DB0-CPB
395425D8          M393A2G40DB0-CPB    M393A2G40DB0-CPB
39542BF2          M393A2G40DB0-CPB    M393A2G40DB0-CPB
395426EF          M393A2G40DB0-CPB    M393A2G40DB0-CPB
3953EB61          M393A2G40DB0-CPB    M393A2G40DB0-CPB
39542AD0          M393A2G40DB0-CPB    M393A2G40DB0-CPB
39542973          M393A2G40DB0-CPB    M393A2G40DB0-CPB
39542ACF          M393A2G40DB0-CPB    M393A2G40DB0-CPB
3953E982          M393A2G40DB0-CPB    M393A2G40DB0-CPB
39542568          M393A2G40DB0-CPB    M393A2G40DB0-CPB
3953EA4A          M393A2G40DB0-CPB    M393A2G40DB0-CPB
3953E993          M393A2G40DB0-CPB    M393A2G40DB0-CPB
804AC8C2          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
707387FA          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
7073895C          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
3077315A          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
90655EDF          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
7073885E          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
30772789          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
707388BB          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
90655FC6          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
70738960          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
70738871          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
7073881A          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
307C97D6          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
7073880D          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
70738819          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
30773136          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
30772FF8          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
70738850          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
3079B2CD          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
307A04D2          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
70738951          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
70738939          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
70738A0B          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
70738823          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
```

Cont:

Output:

```
Requires Service   Service Reason
-----------------  ------------------
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
hardware_status	Hardware Status	4
dimm_id	DIMM Id	N/A
cpu	CPU	N/A
size	Size(Mb)	N/A
speed	Speed(MHz)	N/A
configured_speed	Configured Clock Speed(MHz)	N/A
manufacturer	Manufacturer	N/A
serial	Serial	N/A
original_serial	Original Serial	N/A
part_number	Part #	N/A
original_part_number	Original Part Number	N/A
requires_service	Requires Service	N/A
service_reason	Service Reason	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing CPUs in the modules

Use the **cpu_list** command to list the central processing units (CPU) in the modules.

```
cpu_list [ module=ModuleNumber  | cpu=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
module	Limits the listing to a specific module.	N	All CPUs in all modules.
cpu	Lists only a specific CPU.	N	A specific CPU.

Example:

```
cpu_list
```

Output:

```
Component ID   Status   Currently Functioning   Hardware Status   CPU Number   Family
-------------- -------- ----------------------- ----------------- ------------ --------
1:CPU:10:1     OK       yes                     OK                1            Xeon
1:CPU:11:1     OK       yes                     OK                1            Xeon
1:CPU:12:1     OK       yes                     OK                1            Xeon
1:CPU:13:1     OK       yes                     OK                1            Xeon
1:CPU:1:1      OK       yes                     OK                1            Xeon
1:CPU:2:1      OK       yes                     OK                1            Xeon
1:CPU:3:1      OK       yes                     OK                1            Xeon
1:CPU:4:1      OK       yes                     OK                1            Xeon
1:CPU:5:1      OK       yes                     OK                1            Xeon
1:CPU:6:1      OK       yes                     OK                1            Xeon
1:CPU:7:1      OK       yes                     OK                1            Xeon
1:CPU:8:1      OK       yes                     OK                1            Xeon
1:CPU:9:1      OK       yes                     OK                1            Xeon
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3

Field ID	Field output	Default position
hardware_status	Hardware Status	4
number	CPU Number	5
family_string	Family	6
type_string	Type	N/A
id	ID	N/A
type	Type Code	N/A
family	Family Code	N/A
model	Model Code	N/A
stepping	Stepping	N/A
max_speed	Max Speed(MHz)	N/A
current_speed	Current Speed(MHz)	N/A
status_string	Status	N/A
manufacturer	Manufacturer	N/A
version	Version	N/A
model_string	Model	N/A
signature	Signature	N/A
core_count	Cores	N/A
core_enabled	Enabled Cores	N/A
thread_count	Threads	N/A
serial	Serial	N/A
original_serial	Original Serial	N/A
part_number	Part #	N/A
original_part_number	Original Part Number	N/A
requires_service	Requires Service	N/A
service_reason	Service Reason	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing InfiniBand HCA Adapters in the storage system

Use the **hca_list** command to list the InfiniBand host card adapters (HCAs) in the storage system.

```
hca_list [ module=ModuleNumber | hca=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
module	Limits the listing to a specific module.	N	All InfiniBand HCA adapters in all modules.
hca	Lists only a specific HCA.	N	A specific InfiniBand HCA.

Example:

```
hca_list
```

Output:

```
Component ID   Status   Currently Functioning   Board Description
-------------- -------- ----------------------- --------------------------
1:HCA:10:1     OK       yes                     CB194A - Connect-IB QSFP
1:HCA:7:1      OK       yes                     CB194A - Connect-IB QSFP
1:HCA:9:1      OK       yes                     CB194A - Connect-IB QSFP

cont:
Board ID         Part Number
---------------- -------------
MT_1210110019    46W0572
MT_1210110019    46W0572
MT_1210110019    46W0572
```

Field ID	Field output	Default position

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
version	Version	N/A
board_description	Board Description	4
original_board_description	Original Board Description	N/A
board_id	Board ID	5
original_board_id	Original Board ID	N/A
board_type	Board Type	N/A
original_board_type	Original Board Type	N/A
serial	Serial	N/A
original_serial	Original Serial	N/A
part_number	Part Number	6
original_part_number	Original Part Number	N/A
hardware_revision	Hardware Revision	N/A
requires_service	Requires Service	N/A
service_reason	Service Reason	N/A
adapter_id	HCA Id	N/A
guid	GUID	N/A
vendor_part_id	Vendor Part ID	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing CNA adapters in the system

Use the **cna_list** command to list CNA adapters in the storage system.

```
cna_list [ module=ModuleNumber | cna=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
module	Limits the listing to a specific module.	N	All CNA adapters in all modules.
cna	Lists only a specific CNA.	N	A specific CNA.

A converged network adapter (CNA) is a single network interface card that contains both a Fibre Channel host bus adapter and a TCP/IP Ethernet NIC. It connects servers to FC-based storage area networks (SANs) and Ethernet-based local area networks (LANs).

Example:

```
cna_list
```

Output:

```
Component ID   Status   Currently Functioning   Board Description
-------------- -------- ----------------------- -----------------------------
1:CNA:10:1     OK       yes                     CX312B - ConnectX-3 Pro SFP+
1:CNA:9:1      OK       yes                     CX312B - ConnectX-3 Pro SFP+

Board ID         Part Number
---------------- --------------
MT_1200111023    MCX312B-XCCT
MT_1200111023    MCX312B-XCCT
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
version	Version	N/A
board_description	Board Description	4

Field ID	Field output	Default position
original_board_description	Original Board Description	N/A
board_id	Board ID	5
original_board_id	Original Board ID	N/A
board_type	Board Type	N/A
original_board_type	Original Board Type	N/A
serial	Serial	N/A
original_serial	Original Serial	N/A
part_number	Part Number	6
original_part_number	Original Part Number	N/A
hardware_revision	Hardware Revision	N/A
requires_service	Requires Service	N/A
service_reason	Service Reason	N/A
adapter_id	HCA Id	N/A
guid	GUID	N/A
vendor_part_id	Vendor Part ID	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing module LEDs in the system

Use the **module_led_list** command to display the module LED state in the system.

```
module_led_list [ module=ModuleNumber ]
```

Parameters

Name	Description	Mandatory	Default
module	Limits the listing to a specific module.	N	All LEDs in all modules.

Example:

```
module_led_list module=1:module:3
```

Output:

```
Module       LED               State   Color
------------ ----------------- ------- -------
1:Module:3   Battery           Off     na
1:Module:3   CPU 1             Off     na
1:Module:3   CPU 2             Off     na
1:Module:3   CPU Mismatch      Off     na
1:Module:3   Check Log         Off     na
1:Module:3   DIMM 1            Off     na
1:Module:3   DIMM 10           Off     na
1:Module:3   DIMM 11           Off     na
1:Module:3   DIMM 12           Off     na
1:Module:3   DIMM 13           Off     na
1:Module:3   DIMM 14           Off     na
1:Module:3   DIMM 15           Off     na
1:Module:3   DIMM 16           Off     na
1:Module:3   DIMM 17           Off     na
1:Module:3   DIMM 18           Off     na
1:Module:3   DIMM 19           Off     na
1:Module:3   DIMM 2            Off     na
1:Module:3   DIMM 20           Off     na
1:Module:3   DIMM 21           Off     na
1:Module:3   DIMM 22           Off     na
1:Module:3   DIMM 23           Off     na
1:Module:3   DIMM 24           Off     na
1:Module:3   DIMM 3            Off     na
1:Module:3   DIMM 4            Off     na
1:Module:3   DIMM 5            Off     na
1:Module:3   DIMM 6            Off     na
1:Module:3   DIMM 7            Off     na
1:Module:3   DIMM 8            Off     na
1:Module:3   DIMM 9            Off     na
```

```
1:Module:3   Fan 1            Off     na
1:Module:3   Fan 2            Off     na
1:Module:3   Fan 3            Off     na
1:Module:3   Fan 4            Off     na
1:Module:3   Fan 5            Off     na
1:Module:3   Fan 6            Off     na
1:Module:3   Fan Riser1       Off     na
1:Module:3   Fan Riser2       Off     na
1:Module:3   Fault            Off     na
1:Module:3   IMM2 Heartbeat   Blink   Green
1:Module:3   Identify         Off     na
1:Module:3   Internal RAID    Off     na
1:Module:3   PCI 1            Off     na
1:Module:3   PCI 2            Off     na
1:Module:3   PCI 3            Off     na
1:Module:3   PCI 4            Off     na
1:Module:3   PCI 5            Off     na
1:Module:3   PCI 6            Off     na
1:Module:3   PCI 7            Off     na
1:Module:3   PCI 8            Off     na
1:Module:3   Power            On      Green
1:Module:3   SysBrd Fault     Off     na
```

Field ID	Field output	Default position
module	Module	1
led	LED	2
state	State	3
color	Color	4
reported	Event Active	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing data disk devices in the system

Use the **disk_list** command to lists special vault device statuses.

```
disk_list [ module=ModuleNumber  | disk=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
module	Limits the listing to a specific module.	N	All data disk devices in all modules.
disk	Data disk devices for which special statuses are to be listed.	N	All data disk devices.

This command lists the statuses of the data disk devices, including:

- Component generic status
- Data disk device capacity
- Model
- Serial

Note:
The temperature values are indicated in Celsius.

Example:

```
disk_list
```

Output:

```
Component ID          Status   Currently Functioning   Capacity   Target Status   Vendor
--------------------- -------- ----------------------- ---------- --------------- --------
1:disk:14:1           OK       yes                     250GB                      LENOVO-X
1:disk:14:2           OK       yes                     250GB                      LENOVO-X
1:disk:3:1            OK       yes                     250GB                      LENOVO-X
1:disk:3:2            OK       yes                     250GB                      LENOVO-X
1:disk:5:1            OK       yes                     250GB                      LENOVO-X
1:disk:5:2            OK       yes                     250GB                      LENOVO-X
1:disk:6:1            OK       yes                     250GB                      LENOVO-X
```

```
1:disk:6:2            OK       yes                    250GB                        LENOVO-X

Cont.:
Model              Serial     Firmware   FRU        Temperature   Encryption State
-----------------  ---------- ---------- ---------  ------------  ------------------
HUSMR1625ASS20E    0PVGJTPA   P4C9       00NA685    22            Ready
HUSMR1625ASS20E    0PVGHN6A   P4C9       00NA685    22            Ready
HUSMR1625ASS20E    0PVJ2PEA   P4C9       00NA685    21            Ready
HUSMR1625ASS20E    0PVJ251A   P4C9       00NA685    21            Ready
HUSMR1625ASS20E    0PVJ99UA   P4C9       00NA685    20            Ready
HUSMR1625ASS20E    0PVJRB5A   P4C9       00NA685    21            Ready
HUSMR1625ASS20E    0PVJS39A   P4C9       00NA685    20            Ready
HUSMR1625ASS20E    0PVJ9RAA   P4C9       00NA685    21            Ready
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
capacity_in_bytes	Capacity	N/A
capacity	Capacity	4
target_status	Target Status	5
vendor	Vendor	6
original_vendor	Original Vendor	N/A
model	Model	7
original_model	Original Model	N/A
serial	Serial	8
original_serial	Original Serial	N/A
firmware	Firmware	9
original_firmware	Original Firmware	N/A
part_number	FRU	10
original_part_number	Original FRU	N/A
temperature	Temperature	11
encryption_state	Encryption State	12
hw_mon_node_id	Hw Node Owner	N/A
requires_service	Requires Service	N/A
service_reason	Service Reason	N/A
revision	Revision	N/A
drive_pn	Drive P/N	N/A
original_drive_pn	Original Drive P/N	N/A
fru_pn	FRU P/N	N/A
original_fru_pn	Original FRU P/N	N/A
desc.bgd_scan	Background Scan	N/A
desc.disk_id	Device ID	N/A
desc.last_sample_serial	Last Sample Serial	N/A
desc.last_sample_time	Last Sample Time	N/A
desc.power_is_on	Power On	N/A
desc.power_on_hours	Power On Hours	N/A
desc.power_on_minutes	Power On Minutes	N/A
desc.last_time_pom_was_mod	Last Time Power On Minutes Was Modified	N/A
desc.read_fail	Read Fail	N/A
desc.smart_code	SMART Code	N/A
desc.smart_fail	SMART Fail	N/A
desc.temperature_status.reported_severity	Reported Temperature Severity	N/A
desc.temperature_status.reported_temperature	Reported Temperature	N/A
desc.temperature_status.temperature	Device Temperature	N/A
desc.sw_encryption_active	Software-Based Encryption Active	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Allowed
Read-only users	Allowed
Technicians	Allowed

Listing service statuses

Use the **service_list** command to list all service specific statuses.

```
service_list [ service=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
service	The service to be listed.	N	All services

This command lists the statuses that apply to services. The list includes the following information:

- Component generic status
- Service on/failed
- Comment (optional)

Example:

```
service_list
```

Output:

```
Component ID          Status   Currently Functioning   Target Status
--------------------- -------- ---------------------- ---------------
1:Data:10             OK       yes
1:Data:11             OK       yes
1:Data:7              OK       yes
1:Data:9              OK       yes
1:Data_Reduction:10   OK       yes
1:Data_Reduction:11   OK       yes
1:Data_Reduction:7    OK       yes
1:Data_Reduction:9    OK       yes
1:Interface:10        OK       yes
1:Interface:11        OK       yes
1:Interface:9         OK       yes
1:Remote:10           OK       yes
1:Remote:11           OK       yes
1:Remote:9            OK       yes
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
target_status	Target Status	4

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing system components that require service

Use the **component_service_required_list** command to list system components and their status.

```
component_service_required_list [ component=ComponentId ] [ filter=<ALL|FAILED|NOTOK> ]
```

Parameters

Name	Type	Description	Mandatory	Default
component	N/A	Lists only this component.	N	All components.
filter	Enumeration	Filters the list to show only failed or only non-OK components.	N	ALL

The list can be filtered to show only a specific component, all failed components, or all components in a non-OK state.

For status and configuration of specific component types, refer to the _list commands for specific components, such as: module_list or switch_list. The output is a list of components with the following information for each component:

- Component identification
- Component general status
- Indication about whether the component is currently functioning

Example:

```
component_service_required_list
```

Output:

```
Component ID            Status   Currently Functioning   Requires Service
----------------------- -------- ----------------------- ------------------
1:IB_Module_Port:12:2   Failed   no                      COMPONENT_TEST
1:IB_Module_Port:13:2   Failed   no                      COMPONENT_TEST
1:IB_Module_Port:8:2    Failed   no                      COMPONENT_TEST
1:IB_Switch_Port:2:12   Failed   no                      COMPONENT_TEST
1:IB_Switch_Port:2:13   Failed   no                      COMPONENT_TEST
1:IB_Switch_Port:2:8    Failed   no                      COMPONENT_TEST
```

```
Cont.:

Service Reason
--------------------------
IB_SWITCH_PHY_PORT_NOT_UP
IB_SWITCH_PHY_PORT_NOT_UP
IB_SWITCH_PHY_PORT_NOT_UP
IB_SWITCH_PHY_PORT_NOT_UP
IB_SWITCH_PHY_PORT_NOT_UP
IB_SWITCH_PHY_PORT_NOT_UP
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
requires_service	Requires Service	4
service_reason	Service Reason	5

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing trace snapshot on a module.

Use the **traces_snapshot_list** command to list trace snapshots on a module.

```
traces_snapshot_list module=ModuleNumber
```

Parameters

Name	Description	Mandatory
module	Component ID of the module to query.	Y

Field ID	Field output	Default position
snapshot	Snapshot Directories	1

Example:

```
traces_snapshot_list  module=1:Module:9
```

Output:

```
Snapshots Directories
--------------------------------
1__20120802_1653_20120802_1713
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Allowed

Creating a trace snapshot

Use the **traces_snapshot** command to create a trace snapshot.

```
traces_snapshot [ snapshot_back_time=MINUTES ] [ snapshot_delay_time=MINUTES ]
```

Parameters

Name	Type	Description	Mandatory	Default
snapshot_delay_time	Integer	Max delay between the request and snapshot creation.	N	no. Uses configuration misc.internal.auto_ snapshot_ trace.last_snapshot_ minutes_delay field.
snapshot_back_time	Integer	Time back from the request time to include in the snapshot.	N	no. Uses configuration misc.internal.auto_ snapshot_ trace.snapshot_ back_time field.

Example:

```
traces_snapshot snapshot_back_time=60 snapshot_delay_time=1
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Allowed

Statistics commands

This section describes the command-line interface (CLI) for getting system statistics.

- **Getting performance statistics**
  Use the **statistics_get** command to retrieve performance statistics from the storage system.
- **Retrieving usage history**
  Use the **usage_get** command to display the usage history of a volume or a storage pool.

Getting performance statistics

Use the **statistics_get** command to retrieve performance statistics from the storage system.

```
statistics_get [ perf_class=perfClassName | host=HostName | host_iscsi_name=initiatorName | host_fc_port=WWPN
| target=RemoteTarget | remote_fc_port=WWPN | remote_ipaddress=IPAddress | vol=VolName | domain=DomainName |
ipinterface=IPInterfaceName | local_fc_port=ComponentId ]
 < start=TimeStamp | end=TimeStamp > [ module=ModuleNumber ] count=N interval=IntervalSize resolution_unit=
<minute|hour|day|week|month>
```

Parameters

Name	Type	Description	Mandatory	Default
host	Object name	Limits statistics to the specific host only.	N	All hosts
host_fc_port	N/A	FC address of the host port.	N	All ports.
target	Object name	Limits statistics to I/O generated by the specified remote target only (due to remote mirroring).	N	All targets.
remote_fc_port	N/A	Limits statistics to the specified host/remote FC port only.	N	All ports.
remote_ipaddress	N/A	IP address of the remote target port.	N	All ports.

Name	Type	Description	Mandatory	Default
host_iscsi_name	iSCSI initiator name	Limits statistics to the specified iSCSI initiator only.	N	All ports.
ipinterface	Object name	Limits statistics to the specified IP interface (relevant for iSCSI only).	N	All interfaces.
module	N/A	Limits statistics to the specified module only.	N	All modules.
local_fc_port	N/A	Limits statistics to I/O performed on the specified FC port only.	N	All ports.
vol	Object name	Limits statistics to the specified volume only.	N	All volumes.
domain	Object name	Limits statistics to the specified domain only.	N	All domains.
start	N/A	Starting point for the statistics report.	N	N/A
end	N/A	Ending point for the statistics report.	N	N/A
count	Positive integer	Number of time points reported.	Y	N/A
interval	Positive integer	The length of time in each statistic's time point. The resolution of this number is set in *resolution_unit*.	Y	N/A
resolution_unit	Enumeration	Sets the unit of measurement for the length of each bin.	Y	N/A
perf_class	Object name	Displays performance class aggregated statistics for bandwidth and IOPS.	N	All Performance classes.

This command lists I/O statistics. The count parameter sets the number of lines in the statistics report. The combination of the interval and resolution_unit parameters sets the length of time for each statistics line. Either start or end timestamp must be provided. These timestamps set the time for the statistics report. Other parameters restrict statistics to a specific host, host port, volume, domain, interface port and so on.

For each line of statistics, 48 numbers are reported, which represent all the combinations of reads/writes, hits/misses and I/O size reporting for each of the 16 options for bandwidth, IOPS and latency. Statistics collection is limited to 32 pools and 200 volumes.

The syntax for the start and end fields is as follows: **`Y-M-D[.[h[:m[:s]]]]`**, where the ranges are as follows:

- Y - year (four digit)
- M - month (1-12)
- D - day (1-31)
- h - hour (0-23, with 0 as default)
- m - minute (0-59, with 0 as default)
- s - second (0-59, with 0 as default)

The year, month and day are separated by dashes, and the optional hours, minutes and seconds are separated by colons.

Output units:

- Very Large blocks are >512KB
- Large blocks - 64-512KB
- Medium blocks - 8-64KB
- Small blocks - 0-8KB
- The latency is in Microseconds
- The bandwidth is in KB

Field ID	Field output	Default position
time	Time	1
failures	Failures	N/A
aborts	Aborts	N/A
read_hit_very_large_iops	Read Hit Very large - IOps	2
read_hit_very_large_latency	Read Hit Very large - Latency	3
read_hit_very_large_internal_latency	Read Hit Very large - Internal Latency	75
read_hit_very_large_throughput	Read Hit Very large - Throughput	4
read_hit_very_large_remotely_served	Read Hit Very large - Remotely Served IOs	63
read_hit_large_iops	Read Hit Large - IOps	5
read_hit_large_latency	Read Hit Large - Latency	6
read_hit_large_internal_latency	Read Hit Large - Internal Latency	76
read_hit_large_throughput	Read Hit Large - Throughput	7
read_hit_large_remotely_served	Read Hit Large - Remotely Served IOs	64
read_hit_medium_iops	Read Hit Medium - IOps	8
read_hit_medium_latency	Read Hit Medium - Latency	9
read_hit_medium_internal_latency	Read Hit Medium - Internal Latency	77
read_hit_medium_throughput	Read Hit Medium - Throughput	10
read_hit_medium_remotely_served	Read Hit Medium - Remotely Served IOs	65
read_hit_small_iops	Read Hit Small - IOps	11
read_hit_small_latency	Read Hit Small - Latency	12
read_hit_small_internal_latency	Read Hit Small - Internal Latency	78
read_hit_small_throughput	Read Hit Small - Throughput	13
read_hit_small_remotely_served	Read Hit Small - Remotely Served IOs	66
read_miss_very_large_iops	Read Miss Very large - IOps	14
read_miss_very_large_latency	Read Miss Very large - Latency	15
read_miss_very_large_internal_latency	Read Miss Very large - Internal Latency	79
read_miss_very_large_throughput	Read Miss Very large - Throughput	16

Field ID	Field output	Default position
read_miss_very_large_remotely_served	Read Miss Very large - Remotely Served IOs	67
read_miss_large_iops	Read Miss Large - IOps	17
read_miss_large_latency	Read Miss Large - Latency	18
read_miss_large_internal_latency	Read Miss Large - Internal Latency	80
read_miss_large_throughput	Read Miss Large - Throughput	19
read_miss_large_remotely_served	Read Miss Large - Remotely Served IOs	68
read_miss_medium_iops	Read Miss Medium - IOps	20
read_miss_medium_latency	Read Miss Medium - Latency	21
read_miss_medium_internal_latency	Read Miss Medium - Internal Latency	81
read_miss_medium_throughput	Read Miss Medium - Throughput	22
read_miss_medium_remotely_served	Read Miss Medium - Remotely Served IOs	69
read_miss_small_iops	Read Miss Small - IOps	23
read_miss_small_latency	Read Miss Small - Latency	24
read_miss_small_internal_latency	Read Miss Small - Internal Latency	82
read_miss_small_throughput	Read Miss Small - Throughput	25
read_miss_small_remotely_served	Read Miss Small - Remotely Served IOs	70
write_hit_very_large_iops	Write Hit Very large - IOps	26
write_hit_very_large_latency	Write Hit Very large - Latency	27
write_hit_very_large_internal_latency	Write Hit Very large - Internal Latency	83
write_hit_very_large_throughput	Write Hit Very large - Throughput	28
write_hit_large_iops	Write Hit Large - IOps	29
write_hit_large_latency	Write Hit Large - Latency	30
write_hit_large_internal_latency	Write Hit Large - Internal Latency	84
write_hit_large_throughput	Write Hit Large - Throughput	31
write_hit_medium_iops	Write Hit Medium - IOps	32
write_hit_medium_latency	Write Hit Medium - Latency	33
write_hit_medium_internal_latency	Write Hit Medium - Internal Latency	85
write_hit_medium_throughput	Write Hit Medium - Throughput	34
write_hit_small_iops	Write Hit Small - IOps	35
write_hit_small_latency	Write Hit Small - Latency	36
write_hit_small_internal_latency	Write Hit Small - Internal Latency	86
write_hit_small_throughput	Write Hit Small - Throughput	37
write_miss_very_large_iops	Write Miss Very large - IOps	38
write_miss_very_large_latency	Write Miss Very large - Latency	39
write_miss_very_large_internal_latency	Write Miss Very large - Internal Latency	87
write_miss_very_large_throughput	Write Miss Very large - Throughput	40
write_miss_large_iops	Write Miss Large - IOps	41
write_miss_large_latency	Write Miss Large - Latency	42
write_miss_large_internal_latency	Write Miss Large - Internal Latency	88
write_miss_large_throughput	Write Miss Large - Throughput	43
write_miss_medium_iops	Write Miss Medium - IOps	44
write_miss_medium_latency	Write Miss Medium - Latency	45
write_miss_medium_internal_latency	Write Miss Medium - Internal Latency	89
write_miss_medium_throughput	Write Miss Medium - Throughput	46
write_miss_small_iops	Write Miss Small - IOps	47
write_miss_small_latency	Write Miss Small - Latency	48
write_miss_small_internal_latency	Write Miss Small - Internal Latency	90
write_miss_small_throughput	Write Miss Small - Throughput	49
read_memory_hit_very_large_iops	Read Memory-Hit Very large - IOps	50
read_memory_hit_very_large_latency	Read Memory-Hit Very large - Latency	51
read_memory_hit_very_large_internal_latency	Read Memory-Hit Very large - Internal Latency	91
read_memory_hit_very_large_throughput	Read Memory-Hit Very large - Throughput	52
read_memory_hit_very_large_remotely_served	Read Memory-Hit Very large - Remotely Served IOs	71
read_memory_hit_large_iops	Read Memory-Hit Large - IOps	53
read_memory_hit_large_latency	Read Memory-Hit Large - Latency	54
read_memory_hit_large_internal_latency	Read Memory-Hit Large - Internal Latency	92
read_memory_hit_large_throughput	Read Memory-Hit Large - Throughput	55
read_memory_hit_large_remotely_served	Read Memory-Hit Large - Remotely Served IOs	72
read_memory_hit_medium_iops	Read Memory-Hit Medium - IOps	56
read_memory_hit_medium_latency	Read Memory-Hit Medium - Latency	57
read_memory_hit_medium_internal_latency	Read Memory-Hit Medium - Internal Latency	93
read_memory_hit_medium_throughput	Read Memory-Hit Medium - Throughput	58
read_memory_hit_medium_remotely_served	Read Memory-Hit Medium - Remotely Served IOs	73
read_memory_hit_small_iops	Read Memory-Hit Small - IOps	59

Field ID	Field output	Default position
read_memory_hit_small_latency	Read Memory-Hit Small - Latency	60
read_memory_hit_small_internal_latency	Read Memory-Hit Small - Internal Latency	94
read_memory_hit_small_throughput	Read Memory-Hit Small - Throughput	61
read_memory_hit_small_remotely_served	Read Memory-Hit Small - Remotely Served IOs	74
time_in_seconds	Time (s)	62

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Return codes

- BAD_TIME_FORMAT
  Bad time format. Should be YYYY-MM-DD[.HH[:MM[:SS]]].
- TARGET_PORT_BAD_ADDRESS
  The remote port address is illegal or does not belong to the remote target.
- VOLUME_BAD_NAME
  The volume name does not exist.
- STATS_TOO_MANY_SAMPLES
  The requested number of statistics samples is too high.
- TARGET_BAD_NAME
  The target name does not exist.
- COMPONENT_DOES_NOT_EXIST
  The component does not exist.
- HOST_BAD_NAME
  The host name does not exist.
- HOST_PORT_DOES_NOT_EXIST
  The port ID is not defined.
- IPINTERFACE_DOES_NOT_EXIST
  This IP interface name does not exist.
- PERF_CLASS_BAD_NAME
  The performance class does not exist.
- COMMAND_AMBIGUOUS
  The user belongs to more than one domain. Please specify a domain or an object.
- DOMAIN_DOESNT_EXIST
  The domain does not exist.

Retrieving usage history

Use the **usage_get** command to display the usage history of a volume or a storage pool.

```
usage_get < vol=VolName | pool=PoolName > [ start=TimeStamp | start_in_seconds=StartTime ] [ end=TimeStamp ]
 [ max=MaxEntries ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	Volume for which usage statistics are retrieved.	N	N/A
pool	Object name	Storage pool for which usage statistics are retrieved.	N	N/A
start	N/A	Starting time for usage history retrieval.	N	Creation time of the object.
end	N/A	Ending time for usage history retrieval.	N	Current time.
max	Integer	Maximum number of entries to retrieve.	N	No limit.

Name	Type	Description	Mandatory	Default
start_in_seconds	Integer	Starting time for usage history retrieval, in seconds since 12:00:00 AM, 1 January 1970.	N	Creation time of the object.

This command retrieves the usage history of a storage pool or volume in megabytes (MB).

Example:

```
usage_get pool=DBPool
```

Output:

```
Time                Volume Usage (MiB)   Snapshot Usage (MiB)
------------------- -------------------- ---------------------
2016-03-29 12:00:00 0                    0
2016-03-29 13:00:00 0                    0
2016-03-29 14:00:00 0                    0
```

Field ID	Field output	Default position
time	Time	1
volume_usage	Volume Usage (MiB)	2
snapshot_usage	Snapshot Usage (MiB)	3

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.
- POOL_DOES_NOT_EXIST
  The storage pool does not exist.
- BAD_TIME_FORMAT
  Bad time format. Should be YYYY-MM-DD[.HH[:MM[:SS]]].
- END_BEFORE_START
  The end time cannot precede the start time.
- VOLUME_IS_SNAPSHOT
  The operation is not permitted on snapshots.

Metadata commands

This section describes the command-line interface (CLI) for handling metadata.

- **Setting metadata**
  Use the **metadata_set** command to set metadata of an object.
- **Deleting metadata**
  Use the **metadata_delete** command to delete an object's metadata.
- **Listing metadata**
  Use the **metadata_list** command to list an object's metadata.
- **Setting user metadata**
  Use the **user_metadata_set** command to set user metadata by section and key.
- **Listing user metadata**
  Use the **user_metadata_list** command to list User meta data.
- **Deleting user metadata**
  Use the **user_metadata_delete** command to delete user metadata by user section and key.

Setting metadata

Use the **metadata_set** command to set metadata of an object.

```
metadata_set object_type=Object name=Name key=Key value=Value
```

Parameters

Name	Type	Description	Mandatory
object_type	Enumeration	An object type. Available values: cg, cluster, dest, destgroup, host, performanceclass, pool, rule, schedule, smsgw, smtpgw, target, user, user_group, vol.	Y
name	Object name	An object name.	Y
key	String	Metadata key.	Y
value	String	Metadata value.	Y

This command sets a new metadata key value for the specified object. The new value overrides the previous one, if it exists.

The value can be an empty string. Up to 16 values are allowed, each limited to 128 bytes.

Example:

```
metadata_set object_type=host name=Host1 key=01 value=Host
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission	Condition
Storage administrator	Allowed	N/A
Storage integration administrator	Allowed	N/A
Application administrator	Conditionally Allowed	Metadata can be set for only volumes, snapshots, snapshot groups, clusters or hosts, and only for objects associated with the application administrator executing the command. Hosts or clusters should be associated with the user. Volumes should be mapped to a host or a cluster associated with the user. Snapshots or snapshot groups should be ones created by application administrator.
Security administrator	Disallowed	N/A
Read-only users	Disallowed	N/A
Technicians	Disallowed	N/A

Return codes

- OBJECT_BAD_NAME
  The referenced object does not exist.
- MAX_METADATA_OBJECTS_REACHED
  The maximum number of metadata objects has been reached.
- REMOTE_MAX_METADATA_OBJECTS_REACHED
  The maximum number of metadata objects has been reached on a remote system.
- LOCAL_PEER_IS_NOT_MASTER
  The local peer is not primary.
- TARGET_NOT_CONNECTED
  There is currently no connection to the target system.
- REMOTE_TARGET_NOT_CONNECTED
  There is currently no connection from the target system.
- HA_IS_NOT_OPERATIONAL
  This HyperSwap relation is not operational. The operation cannot be carried out on a non-operational HyperSwap relation.
- MAX_METADATA_VOL_OBJECTS_REACHED
  The maximum number of metadata objects for a volume has been reached.

Deleting metadata

Use the **metadata_delete** command to delete an object's metadata.

```
metadata_delete object_type=Object name=Name key=Key
```

Parameters

Name	Type	Description	Mandatory
object_type	Enumeration	Type of object. Available values: cg, cluster, dest, destgroup, host, performanceclass, pool, rule, schedule, smsgw, smtpgw, target, user, user_group, vol.	Y
name	Object name	The name of the target object.	Y
key	String	Metadata key.	Y

This command deletes a metadata key value for the specified object.

The command will fail if the key is not defined.

Example:

```
metadata_delete object_type=host name=Host1 key=01
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission	Condition
Storage administrator	Allowed	N/A
Storage integration administrator	Allowed	N/A
Application administrator	Conditionally Allowed	Metadata can be set for only volumes, snapshots, snapshot groups, clusters or hosts, and only for objects associated with the application administrator executing the command. Hosts or clusters should be associated with the user. Volumes should be mapped to a host or a cluster associated with the user. Snapshots or snapshot groups should be ones created by application administrator.
Security administrator	Disallowed	N/A
Read-only users	Disallowed	N/A
Technicians	Disallowed	N/A

Return codes

- OBJECT_BAD_NAME
  The referenced object does not exist.
- METADATA_OBJECT_KEY_NOT_FOUND
  The specified metadata object does not exist.
- LOCAL_PEER_IS_NOT_MASTER
  The local peer is not primary.
- HA_IS_NOT_OPERATIONAL
  This HyperSwap relation is not operational. The operation cannot be carried out on a non-operational HyperSwap relation.
- TARGET_NOT_CONNECTED
  There is currently no connection to the target system.
- REMOTE_TARGET_NOT_CONNECTED
  There is currently no connection from the target system.

Listing metadata

Use the **metadata_list** command to list an object's metadata.

```
metadata_list [ object_type=Object ] [ name=Name ] [ key=Key ] [ domain=DomainName ]
```

Parameters

Name	Type	Description	Mandatory	Default
object_type	Enumeration	Type of object.	N	Type of object. Available values: cg, cluster, dest, destgroup, host, performanceclass, pool, rule, schedule, smsgw, smtpgw, target, user, user_group, vol.
name	Object name	The name of the target object.	N	All objects
key	String	Metadata key.	N	List all keys and values.
domain	Object name	The domain name.	N	All Domains

This command lists all the value key pairs for this object, or a specific one. The command fails if no key is defined.

Example:

```
metadata_list object_type=host
```

Output:

```
Object Type    Name     Key    Value
-------------- -------- ------ --------
Host           Host1    01     Host
```

Field ID	Field output	Default position
object_type	Object Type	1
name	Name	2
key	Key	3
value	Value	4

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Setting user metadata

Use the **user_metadata_set** command to set user metadata by section and key.

```
user_metadata_set  section=Section key=Key value=Value
```

Parameters

Name	Type	Description	Mandatory
section	Enumeration	Metadata section.	Y
key	String	Metadata key.	Y
value	String	Metadata value.	Y

Example:

```
user_metadata_set section=GUI key=GUI_10.2 value=Ubunto_2.6
```

Output:

```
(Command returns no output)
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- USER_METADATA_FULL
  Maximal number of user metadata objects has been reached.

Listing user metadata

Use the **user_metadata_list** command to list User meta data.

```
user_metadata_list [ section=Section ]
```

Parameters

Name	Type	Description	Mandatory	Default
section	Enumeration	User metadata section.	N	''

Lists User data.

Example:

```
user_metadata_list
```

Output:

```
Time                 Section  Key        Value
-------------------- -------- ---------- -------
2004-11-22 18:08:23   TA       GUI_0.2    Ubunto_2.2
2007-11-22 18:08:22   GUI      GUI_10.1   Ubunto_2.6
2007-11-22 18:08:23   GUI      GUI_10.2   Ubunto_2.6
```

Field ID	Field output	Default position
time	Time	1
section	Section	2
key	Key	3
value	Value	4

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Deleting user metadata

Use the **user_metadata_delete** command to delete user metadata by user section and key.

```
user_metadata_delete section=Section key=Key
```

Parameters

Name	Type	Description	Mandatory
section	Enumeration	Meta data section.	Y
key	String	Metadata key.	Y

Example:

```
user_metadata_delete section=GUI key=GUI_10.2
```

Output:

```
(Command returns no output)
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- USER_METADATA_OBJECT_KEY_NOT_FOUND
  The specified user metadata object does not exist.

Encryption enablement and support commands

This section describes the command-line interface (CLI) for encryption configuration.

- **Disabling encryption**
  Use the **encrypt_disable** command to disable the data protection feature.
- **Enabling encryption**
  Use the **encrypt_enable** command to enable the data protection feature.
- **Defining a key server**
  Use the **encrypt_keyserver_define** command to define a key server to be used by the system.
- **Removing a key server**
  Use the **encrypt_keyserver_delete** command to remove a key server used by the system.
- **Displaying key server status**
  Use the **encrypt_keyserver_list** command to list the key servers currently defined in the system along with their connectivity status.
- **Checking key server status**
  Use the **encrypt_keyserver_check_status** command to check connectivity status of the key servers currently defined in the system.
- **Obtaining a new master key**
  Use the **encrypt_keyserver_rekey** command to initiate a rekey against the master key server.
- **Renaming a key server**
  Use the **encrypt_keyserver_rename** command to change the name of a defined key server.
- **Changing key server properties**
  Use the **encrypt_keyserver_update** command to change a key server's IP address and/or port.
- **Entering a recovery key**
  Use the **encrypt_recovery_key_enter** command to unlock encrypted disks when the system reboots and cannot access any of the defined keyservers, and when recovery keys are defined.
- **Generating recovery keys**
  Use the **encrypt_recovery_key_generate** command to specify which security administrators will receive recovery key shares, and to define the minimum number of recovery key shares that need to be entered.
- **Retrieving the security administrator's recovery key**
  Use the **encrypt_recovery_key_get** command to retrieve the recovery key share generated for the current user.
- **Rekeying the security administrators**
  Use the **encrypt_recovery_key_rekey** command to restart the recovery key generation process.
- **Displaying recovery key status**
  Use the **encrypt_recovery_key_status** command to display status information for recovery keys.
- **Recovering key verification**
  Use the **encrypt_recovery_key_verify** command to confirm that the current user has correctly copied the recovery key share retrieved by the **encrypt_recovery_key_get** command.
- **Recovering key share information**
  Use the **encrypt_recovery_key_list** command to list recovery key share information.
- **Finishing the recovery process**
  Use the **encrypt_recovery_finish** command to finish the recovery process and move the system to the On state.
- **Obtaining a new master key**
  Use the **encrypt_local_rekey** command to initiate rekeying (getting new cryptographic material) from the local key management.
- **Changing the key management scheme**
  Use the **encrypt_change_key_scheme** command to change the key management scheme.
- **Viewing the key scheme**
  Use the **encrypt_key_scheme_get** command to view the key scheme defined in the system.

Disabling encryption

Use the **encrypt_disable** command to disable the data protection feature.

`encrypt_disable`

This command disables the data protection feature.

In order for this command to complete successfully, all of the following prerequisites must be fulfilled:

- The system is fully redundant
- None of the present and active SSDs, Flash cards, or Flash canisters have failed
- No Flash enclosure is undergoing CCL

When data protection is disabled, a cryptographic erase is performed on all protected bands (ensuring that all existing user data is no longer accessible). After the command successfully completes, all bands are left in the unlocked state.

Disabling encryption when the encryption state is other than Active (displayed as **`Enabled`** in state_list) will result in an error.

Example:

```
encrypt_disable -y
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Disallowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_DISABLE_ENCRYPTION
  Are you sure you want to disable encryption on this system?

  Troubleshooting: To proceed with issuing the command, enter -y.

Return codes

- CANNOT_DISABLE_ENCRYPTION_WHILE_NOT_IN_FULLY_PROTECTED_STATE
  Cannot disable encryption while not in a fully protected state.

  Troubleshooting: Resolve any issues preventing system from reaching full data protection state. Contact IBM Support.

- UNSUPPORTED_HARDWARE
  Cannot encrypt on unsupported hardware.

  Troubleshooting: Contact IBM support to verify encryption status.

- ENCRYPT_NOT_ENABLED
  Encryption is not enabled.

  Troubleshooting: Make sure that encryption is enabled and re-run the command.

- VOLUME(S)_DEFINED
  There are volumes defined, cannot disable encryption.

  Troubleshooting: All volumes must be removed before encryption is disabled.

- CANNOT_UNMOUNT_STATISTIC_VOLUME
  Failed to unmount the statistics volume for disabling encryption.

  Troubleshooting: Contact IBM Support.

- DATA_REDUCTION_RECOVERY_IS_RUNNING
  Data reduction recovery is running, the operation is not allowed.

  Troubleshooting: Wait for data reduction recovery to complete.

- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support

- SECURE_ERASE_IS_RUNNING
  The operation not allowed while secure erase is running.

  Troubleshooting: Wait for the secure erase process to complete.

- FLASH_CCL_IN_PROGRESS
  The requested command cannot be invoked while Flash Enclosure CCL is in progress.

  Troubleshooting: Wait for Flash CCL to complete.

- ENCRYPTION_CANNOT_UNENROLL_SOME_VAULT_DEVICES
  Some vault devices cannot be un-enrolled due to failed components.

  Troubleshooting: Contact IBM Support.

- ENCRYPTION_CANNOT_UNENROLL_SOME_FLASH_ENCLOSURES
  Not all flash enclosures are phased in before un-enrollment.

  Troubleshooting: Contact IBM Support.

- ENCRYPTION_IS_NOT_IN_A_STABLE_STATE
  Encryption is not in a stable state.

  Troubleshooting: Wait for the encryption process to complete.

Enabling encryption

Use the **encrypt_enable** command to enable the data protection feature.

```
encrypt_enable [ recovery_keys=<yes|no> ] [ key_scheme=KeyScheme ]
```

Parameters

Name	Type	Description	Mandatory	Default
recovery_keys	Boolean	Defines whether recovery keys are required for encryption activation.	N	yes
key_scheme	Enumeration	Defines which key scheme to use for encryption activation: external or local.	N	external

This command enables the data protection feature. In order for this command to complete successfully, all of the following prerequisites must be fulfilled:

- The system is fully redundant
- None of the present and active SSDs, Flash cards, or Flash canisters have failed
- No Flash enclosure is undergoing CCL

Example:

```
encrypt_enable recovery_keys=yes key_scheme=local -y
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Disallowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_ENABLE_ENCRYPTION
  Are you sure you want to enable encryption on this system?

  Troubleshooting: To proceed with issuing the command, enter -y.

- ARE_YOU_SURE_YOU_WANT_TO_ENABLE_LOCAL_KEY_MANAGEMENT_ENCRYPTION
  Are you sure you want to enable encryption with local key management on this system?

  Troubleshooting: To proceed with issuing the command, enter -y.

Return codes

- CANNOT_ENABLE_ENCRYPTION_WHILE_NOT_IN_FULLY_PROTECTED_STATE
  Cannot enable encryption while not in a fully protected state.

  Troubleshooting: Resolve any issues preventing the system from reaching a full data protection state, and contact IBM Support.

- UNSUPPORTED_HARDWARE
  Cannot encrypt on unsupported hardware.

  Troubleshooting: Contact IBM support to verify encryption status.

- INVALID_RECOVERY_KEY_STATE
  The recovery key state is inconsistent with the specified option.

  Troubleshooting: Check the recovery key state using the encrypt_recovery_key_status command.

- ENCRYPTION_ALREADY_ENABLED
  Encryption has already been enabled.

  Troubleshooting: Check the output of the state_list command.

- SYSTEM_IS_REDISTRIBUTING

The operation is not allowed during rebuild or phase-in.

- NO_LIVE_KEYSERVER_GATEWAY_NODE
  There is no live key server gateway node in the system.

  Troubleshooting: Restart the key server gateway node and try again.

- NO_MASTER_KEYSERVER_DEFINED
  No master key server is defined in the system.

  Troubleshooting: Define a master key server by invoking encrypt_key server_update and try again.

- KEYSERVER_COMMUNICATION_GENERIC_ERROR
  Cannot connect to an active key server.

  Troubleshooting: Invoke encrypt_keyserver_list and event_list for more details.

- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support

- ENCRYPTION_CANNOT_ENROLL_SOME_FLASH_ENCLOSURES
  Not all flash enclosures are phased in before enrollment.

  Troubleshooting: Contact IBM Support.

- ENCRYPTION_CANNOT_ENROLL_SOME_VAULT_DEVICES
  Some vault devices cannot be enrolled due to failed components.

  Troubleshooting: Contact IBM Support.

- DATA_REDUCTION_RECOVERY_IS_RUNNING
  Data reduction recovery is running, the operation is not allowed.

  Troubleshooting: Wait for data reduction recovery to complete.

- FLASH_CCL_IN_PROGRESS
  The requested command cannot be invoked while Flash Enclosure CCL is in progress.

  Troubleshooting: Wait for Flash CCL to complete.

- RECOVERY_MISMATCH_KEY_SCHEME
  The recovery key scheme does not match the current scheme.

  Troubleshooting: Check the recovery key scheme using the encrypt_key_scheme_get command.

Defining a key server

Use the **encrypt_keyserver_define** command to define a key server to be used by the system.

```
encrypt_keyserver_define name=Name [ ipv4=Address ] [ ipv6=Address ] [ port=PortNumber ] [ master=<yes|no> ] [
keyserver_type=KeyserverType ] certificate=PemCertificate
```

Parameters

Name	Type	Description	Mandatory	Default
name	String	The name of the key server being added.	Y	N/A
certificate	N/A	The public certificate or certificate chain of the key server being added (see below for details).	Y	N/A
master	Boolean	Defines whether this key server is the primary key server used for key retrieval.	N	no
ipv4	N/A	The IPv4 address of the key server being added. Either one IPv4 and/or one IPv6 must be used.	N	NONE
ipv6	N/A	The IPv6 address of the key server being added. Either one IPv4 and/or one IPv6 must be used.	N	NONE
port	Integer	Port used for key server communication.	N	5696
keyserver_type	Enumeration	The type of the key server to communicate with. • Use **'TKLM'** for IBM TKLM, SKLM and iSKLM keyservers • Use **'KeySecure'** for SafeNet, Gemalto KeySecure, Thales KeySecure, and Thales CipherTrust keyservers.	N	TKLM

This command defines a key server to be used by the system upon startup or encryption activation to retrieve the key material required to cryptographically unlock the disks. At least one key server (but preferably two, and no more than four) must be defined and accessible in order for encrypt_enable to succeed. Only one of the key servers may be defined as master.

The certificate parameter:

The value of the certificate parameter is the content of a PEM file with asterisks instead of newlines. Chained certificates are supported. The total maximal length of a PEM file holding chained certificates (leaf first, root last) is 15360 characters (including the asterisk characters). In Windows, you can drag-and-drop a PEM file from the Windows Explorer to the appropriate location in the CLI session window; the content will be added automatically.

Example:

```
encrypt_keyserver_define name=snocone ipv4=snocone.example.com ipv6=2001:ODB8::a5a7 certificate=
"----BEGIN CERTIFICATE---*MIICyTCCAbGgAwIBAgIXLSiyd2FPMA0GCSqGSIb3IiEBCwUAMBQxEjAQAgNVBVuTCXNrbG5pdHNv*
......
*erD5HgQHSkfR3FEM+b6EBOUPFIBrys8rKtLRbWvovobq*---END CERTIFICATE----"
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Disallowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- UNSUPPORTED_HARDWARE
  Cannot encrypt on unsupported hardware.

  Troubleshooting: Contact IBM support to verify encryption status.

- ENCRYPTION_TOO_MANY_KEYSERVERS
  Too many key servers are already defined, cannot add another one.

  Troubleshooting: Delete a key server and try again.

- ENCRYPTION_KEYSERVER_NAME_EXISTS
  The key server name already exists.

  Troubleshooting: Check the currently defined key servers.

- ENCRYPTION_KEYSERVER_MUST_HAVE_ADDRESS
  A key server must have at least one address (IPv4/IPv6).

  Troubleshooting: Make sure the command includes the ipv4= or ipv6= argument.

- ENCRYPTION_KEYSERVER_IPV4_ALREADY_EXISTS
  The IPv4 address or host name already exists.

  Troubleshooting: Check the currently defined key servers.

- ENCRYPTION_KEYSERVER_IPV6_ALREADY_EXISTS
  The IPv6 address or host name already exists.

  Troubleshooting: Check the currently defined key servers.

- SSL_CERTIFICATE_HAS_EXPIRED
  The SSL certificate has expired.

- SSL_CERTIFICATE_VERIFICATION_FAILED
  The SSL certificate chain verification failed.

- SSL_CERTIFICATE_INVALID_FORMAT
  The SSL certificate format is invalid or corrupted.

- SSL_CERTIFICATE_NOT_YET_VALID
  The SSL certificate is not yet valid.

- SSL_CERTIFICATE_VERIFICATION_INTERNAL_ERROR
  The SSL certificate verification has failed because of an internal system error.

- SSL_CERTIFICATE_ISSUER_NOT_FOUND
  The SSL certificate issuer was not found in the certificate chain.

- SSL_CERTIFICATE_CHAIN_EMPTY
  No certificates were found in the input.

Removing a key server

Use the **encrypt_keyserver_delete** command to remove a key server used by the system.

```
encrypt_keyserver_delete name=Name
```

Parameters

Name	Type	Description	Mandatory
name	String	The name of a defined key server.	Y

Example:

```
encrypt_keyserver_delete name=snocone
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Disallowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- UNSUPPORTED_HARDWARE
  Cannot encrypt on unsupported hardware.

  Troubleshooting: Contact IBM support to verify encryption status.

- ENCRYPTION_UNKNOWN_KEYSERVER
  Unknown key server name.

  Troubleshooting: Check the currently defined key servers.

- ENCRYPTION_DELETE_MASTER_KEYSERVER
  Removal of the master key server is not permitted.

  Troubleshooting: A new master key server must be defined before removing the current master.

- ENCRYPTION_LAST_DEFINED_KEYSERVER
  Cannot delete the last key server.

  Troubleshooting: Define another master key server before attempting to delete this one.

Displaying key server status

Use the **encrypt_keyserver_list** command to list the key servers currently defined in the system along with their connectivity status.

```
encrypt_keyserver_list
```

Example:

```
encrypt_keyserver_list
```

Output:

```
Module   Name      App/Key Status   Last time checked   Master   Port
3        nachos    NOAPP            2013/03/27 20:18:43 yes      5696
3        nachos    UNKNOWN          2013/03/27 20:18:43 yes      5696
3        snocone   UNKNOWN          2013/03/27 20:18:43 no       5696
3        snocone   ACTIVE           2013/03/27 20:18:43 no       5696
3        TKLM-SA   BAD_CERT         2013/03/27 20:18:43 no       5696

Address
192.0.2.1
2001:0DB8::1
snocone.example.com
2001:0DB8:e006:238:209:6bff:fe00:a5a7
tklm-sa.example.com
```

Field ID	Field output	Default position
module_id	Module	1

Field ID	Field output	Default position
label	Name	2
heartbeat_keyserver_status	App/Key Status	3
last_heartbeat	Last time checked	4
master	Master	5
port	Port	6
address	Address	7
keyserver_type	Keyserver Type	8

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Allowed
Technicians	Allowed

Return codes

- UNSUPPORTED_HARDWARE
  Cannot encrypt on unsupported hardware.

  Troubleshooting: Contact IBM support to verify encryption status.

Checking key server status

Use the **encrypt_keyserver_check_status** command to check connectivity status of the key servers currently defined in the system.

```
encrypt_keyserver_check_status
```

This command initiates the async check of the connectivity status of the key servers currently defined in the system. To get the current status, use the encrypt_keyserver_list CLI command.

Example:

```
encrypt_keyserver_check_status
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Allowed
Technicians	Allowed

Return codes

- UNSUPPORTED_HARDWARE
  Cannot encrypt on unsupported hardware.

  Troubleshooting: Contact IBM support to verify encryption status.

- CERTIFICATE_FOR_XIV_IS_NOT_INSTALLED
  An XIV encryption certificate is not installed.

  Troubleshooting: Check the output of pki_list for a certificate named XIV and contact IBM Support.

- ENCRYPTION_NOT_IN_EXTERNAL_SCHEME
  Encryption key management is not set to the external scheme.

  Troubleshooting: Check the output of the encrypt_key_scheme_get command.

Obtaining a new master key

Use the **encrypt_keyserver_rekey** command to initiate a rekey against the master key server.

```
encrypt_keyserver_rekey
```

This command initiates a rekeying (getting new cryptographic material) with the master key server.

Example:

```
encrypt_keyserver_rekey
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Disallowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- UNSUPPORTED_HARDWARE
  Cannot encrypt on unsupported hardware.

  Troubleshooting: Contact IBM support to verify encryption status.

- ENCRYPT_NOT_ENABLED
  Encryption is not enabled.

  Troubleshooting: Make sure that encryption is enabled and re-run the command.

- CANNOT_GET_XIV_MASTER_KEY
  Problem obtaining XIV master key from the key server.

  Troubleshooting: Make sure that the key service is active and serving keys (obtained from the key server or stored locally). Otherwise, contact IBM Support.

- CANNOT_GET_NEW_KEY_REQUEST
  Error requesting encryption keys from the key server gateway node.

  Troubleshooting: Make sure that the key server is actively serving keys.

- CANNOT_UPDATE_KEY_METADATA
  Cannot update metadata in the key repository for a new key.

  Troubleshooting: Contact IBM Support.

- CANNOT_CANNOT_GENERATE_EXMK_ESKH
  Problem generating EXMK and ESKH.

  Troubleshooting: Contact IBM Support.

- CANNOT_WRITE_TO_KEY_REPOSITORY
  Failed writing keys to the key repository.

  Troubleshooting: Contact IBM Support.

- CANNOT_COPY_KEYS_IN_KEY_REPOSITORY
  Problem copying current keys to the location of the old keys in the key repository.

  Troubleshooting: Contact IBM Support

- ENCRYPTION_KR_WRITE_FAILED
  Error writing to the key repository.

  Troubleshooting: Contact IBM Support.

- ENCRYPTION_KR_READ_FAILED
  Error reading the key repository.

  Troubleshooting: Contact IBM Support.

- NO_LIVE_KEYSERVER_GATEWAY_NODE

There is no live key server gateway node in the system.

Troubleshooting: Restart the key server gateway node and try again.

- NO_MASTER_KEYSERVER_DEFINED
  No master key server is defined in the system.

  Troubleshooting: Define a master key server by invoking encrypt_key server_update and try again.
- KEYSERVER_COMMUNICATION_GENERIC_ERROR
  Cannot connect to an active key server.

  Troubleshooting: Invoke encrypt_keyserver_list and event_list for more details.
- ENCRYPTION_NOT_IN_EXTERNAL_SCHEME
  Encryption key management is not set to the external scheme.

  Troubleshooting: Check the output of the encrypt_key_scheme_get command.

Renaming a key server

Use the **encrypt_keyserver_rename** command to change the name of a defined key server.

```
encrypt_keyserver_rename name=Name new_name=Name
```

Parameters

Name	Type	Description	Mandatory
new_name	String	The new name of the key server.	Y
name	String	The current name of a defined key server.	Y

Example:

```
encrypt_keyserver_rename name=nachos new_name=snocone
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Disallowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- UNSUPPORTED_HARDWARE
  Cannot encrypt on unsupported hardware.

  Troubleshooting: Contact IBM support to verify encryption status.
- ENCRYPTION_UNKNOWN_KEYSERVER
  Unknown key server name.

  Troubleshooting: Check the currently defined key servers.
- ENCRYPTION_KEYSERVER_NAME_EXISTS
  The key server name already exists.

  Troubleshooting: Check the currently defined key servers.

Changing key server properties

Use the **encrypt_keyserver_update** command to change a key server's IP address and/or port.

```
encrypt_keyserver_update name=Name [ ipv4=Address ] [ ipv6=Address ] [ port=PortNumber ] [ master=<yes|no> ]
 [ certificate=PemCertificate ]
```

Parameters

Name	Type	Description	Mandatory	Default
name	String	Name of the key server to be updated.	Y	N/A
certificate	N/A	The public certificate or certificate chain of the key server to be updated. The certificate parameter: The value of the certificate parameter is the content of a PEM file with asterisks instead of newlines. Chained certificates are supported. The total maximal length of a PEM file holding chained certificates (leaf first, root last) is 15360 characters (including the asterisk characters). In Windows, you can drag-and-drop a PEM file from the Windows Explorer to the appropriate location in the XCLI session window; the content will be added automatically.	N	none
master	Enumeration	Indicates whether this key server is the master.	N	no
ipv4	N/A	The IPv4 address.	N	none
ipv6	N/A	The IPv6 address.	N	none
port	Integer	Port number for communications.	N	5696

This command is used to update a key server's address, port, or certificate.

Example:

```
encrypt_keyserver_update name=nachos master=yes ipv4=192.0.2.1 ipv6=2001:ODB8::2 port=1010 certificate=''
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Disallowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- UNSUPPORTED_HARDWARE
  Cannot encrypt on unsupported hardware.

  Troubleshooting: Contact IBM support to verify encryption status.
- ENCRYPTION_UNKNOWN_KEYSERVER
  Unknown key server name.

  Troubleshooting: Check the currently defined key servers.
- ENCRYPTION_KEYSERVER_IPV4_ALREADY_EXISTS
  The IPv4 address or host name already exists.

  Troubleshooting: Check the currently defined key servers.
- ENCRYPTION_KEYSERVER_IPV6_ALREADY_EXISTS
  The IPv6 address or host name already exists.

  Troubleshooting: Check the currently defined key servers.
- SSL_CERTIFICATE_HAS_EXPIRED
  The SSL certificate has expired.
- SSL_CERTIFICATE_VERIFICATION_FAILED
  The SSL certificate chain verification failed.
- SSL_CERTIFICATE_INVALID_FORMAT
  The SSL certificate format is invalid or corrupted.
- SSL_CERTIFICATE_NOT_YET_VALID
  The SSL certificate is not yet valid.
- SSL_CERTIFICATE_VERIFICATION_INTERNAL_ERROR
  The SSL certificate verification has failed because of an internal system error.
- SSL_CERTIFICATE_ISSUER_NOT_FOUND
  The SSL certificate issuer was not found in the certificate chain.

- SSL_CERTIFICATE_CHAIN_EMPTY
  No certificates were found in the input.

Entering a recovery key

Use the **encrypt_recovery_key_enter** command to unlock encrypted disks when the system reboots and cannot access any of the defined keyservers, and when recovery keys are defined.

```
encrypt_recovery_key_enter key=Key
```

Parameters

Name	Description	Mandatory
key	The 64-character hexadecimal recovery key.	Y

This command is used to unlock encrypted disks when the system reboots and cannot access any of the defined keyservers. To unlock the disks, the min_req number (defined by the encrypt_recovery_key_generate command) of security administrators must all successfully enter their recovery key (as presented to them via recovery_key_get). After the minimum required keys have been entered, the storage administrator must change the state from Maintenance to On by issuing state_change target_state=on. When this command is issued with the machine in the On state, it has no effect, and can be used to check the validity of the recovery key.

Example:

```
encrypt_recovery_key_enter key=CBC9B398373FDE79CD38B23192DABACADB5DA63A915CBF5CA8C4E0C212819DE6
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Disallowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- UNSUPPORTED_HARDWARE
  Cannot encrypt on unsupported hardware.

  Troubleshooting: Contact IBM support to verify encryption status.

- INVALID_RECOVERY_KEY_FRAGMENT
  The provided recovery key fragment does not match the stored key.

  Troubleshooting: Make sure that the proper key (share) has been used.

- GENERIC_FAILED
  Generic encryption failure.

  Troubleshooting: Contact IBM Support.

- INVALID_RECOVERY_KEY_USER
  The user is not a valid recovery key administrator.

  Troubleshooting: Make sure that the provided usernames are valid.

- NO_LIVE_KEYSERVER_GATEWAY_NODE
  There is no live key server gateway node in the system.

  Troubleshooting: Restart the key server gateway node and try again.

- CANNOT_READ_FROM_KEY_REPOSITORY
  Failed reading keys from the key repository.

  Troubleshooting: Contact IBM Support.

- RK_FAILED_VERIFY_SLEEP
  Too many failed verify attempts.

  Troubleshooting: Wait a little and try again.

- ENCRYPTION_KR_WRITE_FAILED

Error writing to the key repository.

Troubleshooting: Contact IBM Support.

- RK_ENTER_SYSTEM_STATE_INVALID
  The command is supported in maintenance mode only.

  Troubleshooting: Switch the system state to maintenance mode.

- INVALID_RECOVERY_KEY_STATE
  The recovery key state is inconsistent with the specified option.

  Troubleshooting: Check the recovery key state using the encrypt_recovery_key_status command.

- RECOVERY_KEY_ALREADY_VERIFIED
  The recovery key has already been verified.

  Troubleshooting: Check the recovery key state using the encrypt_recovery_key_list command.

Generating recovery keys

Use the **encrypt_recovery_key_generate** command to specify which security administrators will receive recovery key shares, and to define the minimum number of recovery key shares that need to be entered.

```
encrypt_recovery_key_generate users=Users [ min_req=MinRequired ] [ key_scheme=KeyScheme ]
```

Parameters

Name	Type	Description	Mandatory	Default
min_req	Integer	Minimum number of required security administrator recovery key shares.	N	2
users	Object name	User names of the security administrators.	Y	N/A
key_scheme	Enumeration	Defines which key scheme to use for encryption activation external or local.	N	external

This command is used to specify which security administrator will receive recovery keys (or, more accurately, "shares"), and to define the minimum number of recovery keys that need to be entered (using the encrypt_recovery_key_enter command) in order to unlock the encrypted keys. Once this command has been entered, all the specified security administrators are expected to retrieve and verify their recovery keys, using encrypt_recovery_key_get and encrypt_recovery_key_verify, respectively. This command can only be run when encryption_state is DISABLED.

Example:

```
encrypt_recovery_key_generate users=secadmin1,secadmin2,secadmin3,secadmin4 min_req=2
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Disallowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- UNSUPPORTED_HARDWARE
  Cannot encrypt on unsupported hardware.

  Troubleshooting: Contact IBM support to verify encryption status.

- NO_LIVE_KEYSERVER_GATEWAY_NODE
  There is no live key server gateway node in the system.

  Troubleshooting: Restart the key server gateway node and try again.

- CANNOT_WRITE_TO_KEY_REPOSITORY
  Failed writing keys to the key repository.

  Troubleshooting: Contact IBM Support.

- CANNOT_GET_NEW_KEY_REQUEST
  Error requesting encryption keys from the key server gateway node.

Troubleshooting: Make sure that the key server is actively serving keys.

- KEYSERVER_COMMUNICATION_GENERIC_ERROR
  Cannot connect to an active key server.

  Troubleshooting: Invoke encrypt_keyserver_list and event_list for more details.

- INSUFFICIENT_RK_ADMIN_THRESHOLD
  Recovery key creation requires at least two security administrators.

  Troubleshooting: Re-run the command by indicating at least two security administrators.

- ENCRYPTION_KR_WRITE_FAILED
  Error writing to the key repository.

  Troubleshooting: Contact IBM Support.

- ENCRYPTION_ALREADY_ENABLED
  Encryption has already been enabled.

  Troubleshooting: Check the output of the state_list command.

- NO_MASTER_KEYSERVER_DEFINED
  No master key server is defined in the system.

  Troubleshooting: Define a master key server by invoking encrypt_key server_update and try again.

- INVALID_RECOVERY_KEY_STATE
  The recovery key state is inconsistent with the specified option.

  Troubleshooting: Check the recovery key state using the encrypt_recovery_key_status command.

- INSUFFICIENT_RK_ADMINS
  The number of users must be greater than or equal to the minimal required number.

  Troubleshooting: Re-run the command by providing at least the minimal number of required users.

- CANNOT_GENERATE_KEYS_ON_KEYSERVER_GATEWAY
  Failed to generate XMK and hashes on a key server gateway node.

  Troubleshooting: Contact IBM Support.

- ENCRYPTION_KR_READ_FAILED
  Error reading the key repository.

  Troubleshooting: Contact IBM Support.

- CANNOT_UPDATE_KEY_METADATA
  Cannot update metadata in the key repository for a new key.

  Troubleshooting: Contact IBM Support.

Retrieving the security administrator's recovery key

Use the **encrypt_recovery_key_get** command to retrieve the recovery key share generated for the current user.

```
encrypt_recovery_key_get
```

This command retrieves the recovery key generated for the current user (by issuing encrypt_recovery_key_generate or encrypt_recovery_key_rekey) to be stored in a secure manner. After running this command, the user needs to 'prove' that they have the key by entering it via the encrypt_recovery_key_verify command. Once this is completed successfully, encrypt_recovery_key_get will no longer present the user's key. Using encrypt_recovery_key_get more than once will return the same value again.

Example:

```
encrypt_recovery_key_get
```

Output:

```
Command completed successfully.
key=B07C4374AC26C4DD3EC2E755EB3FAAF04EC792C8BE0D0CB1C1BAC79998EBEC6D
```

Access control

User Category	Permission
Storage administrator	Disallowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed

User Category	Permission
Technicians	Disallowed

Return codes

- UNSUPPORTED_HARDWARE
  Cannot encrypt on unsupported hardware.

  Troubleshooting: Contact IBM support to verify encryption status.
- INVALID_RECOVERY_KEY_USER
  The user is not a valid recovery key administrator.

  Troubleshooting: Make sure that the provided usernames are valid.
- NO_LIVE_KEYSERVER_GATEWAY_NODE
  There is no live key server gateway node in the system.

  Troubleshooting: Restart the key server gateway node and try again.
- CANNOT_READ_FROM_KEY_REPOSITORY
  Failed reading keys from the key repository.

  Troubleshooting: Contact IBM Support.
- CANNOT_GET_NEW_KEY_REQUEST
  Error requesting encryption keys from the key server gateway node.

  Troubleshooting: Make sure that the key server is actively serving keys.
- KEYSERVER_COMMUNICATION_GENERIC_ERROR
  Cannot connect to an active key server.

  Troubleshooting: Invoke encrypt_keyserver_list and event_list for more details.
- NO_MASTER_KEYSERVER_DEFINED
  No master key server is defined in the system.

  Troubleshooting: Define a master key server by invoking encrypt_key server_update and try again.
- INVALID_RECOVERY_KEY_STATE
  The recovery key state is inconsistent with the specified option.

  Troubleshooting: Check the recovery key state using the encrypt_recovery_key_status command.
- RECOVERY_KEY_ALREADY_VERIFIED
  The recovery key has already been verified.

  Troubleshooting: Check the recovery key state using the encrypt_recovery_key_list command.

Rekeying the security administrators

Use the **encrypt_recovery_key_rekey** command to restart the recovery key generation process.

```
encrypt_recovery_key_rekey [ users=Users ] [ min_req=MinRequired ]
```

Parameters

Name	Type	Description	Mandatory	Default
min_req	Integer	Minimum number of required security administrator recovery key shares.	N	0
users	Object name	Comma delimited list of security administrator to rekey.	N	N/A

This command restarts the recovery key generation process, described in the section on the encrypt_recovery_key_generate command. The only difference is that the parameters users and min_required are optional, and will default to the values specified in the last call to encrypt_recovery_key_generate. Note that none of the new recovery keys will take effect until the last user has verified his or her recovery key. Until then, if recovery is required, the previous keys will remain valid.

Example:

```
encrypt_recovery_key_rekey users=secadmin1,secadmin2,secadmin3,secadmin4 min_req=3
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Disallowed

User Category	Permission
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- UNSUPPORTED_HARDWARE
  Cannot encrypt on unsupported hardware.

  Troubleshooting: Contact IBM support to verify encryption status.

- NO_LIVE_KEYSERVER_GATEWAY_NODE
  There is no live key server gateway node in the system.

  Troubleshooting: Restart the key server gateway node and try again.

- CANNOT_READ_FROM_KEY_REPOSITORY
  Failed reading keys from the key repository.

  Troubleshooting: Contact IBM Support.

- INSUFFICIENT_RK_ADMIN_THRESHOLD
  Recovery key creation requires at least two security administrators.

  Troubleshooting: Re-run the command by indicating at least two security administrators.

- ENCRYPTION_KR_WRITE_FAILED
  Error writing to the key repository.

  Troubleshooting: Contact IBM Support.

- NO_MASTER_KEYSERVER_DEFINED
  No master key server is defined in the system.

  Troubleshooting: Define a master key server by invoking encrypt_key server_update and try again.

- INVALID_RECOVERY_KEY_STATE
  The recovery key state is inconsistent with the specified option.

  Troubleshooting: Check the recovery key state using the encrypt_recovery_key_status command.

- INSUFFICIENT_RK_ADMINS
  The number of users must be greater than or equal to the minimal required number.

  Troubleshooting: Re-run the command by providing at least the minimal number of required users.

- CANNOT_GENERATE_KEYS_ON_KEYSERVER_GATEWAY
  Failed to generate XMK and hashes on a key server gateway node.

  Troubleshooting: Contact IBM Support.

- KEYSERVER_COMMUNICATION_GENERIC_ERROR
  Cannot connect to an active key server.

  Troubleshooting: Invoke encrypt_keyserver_list and event_list for more details.

Displaying recovery key status

Use the **encrypt_recovery_key_status** command to display status information for recovery keys.

```
encrypt_recovery_key_status
```

This command shows status information regarding recovery keys, specifically: Which user has verified his or her recovery key before encrypt_enable or in the recovery key rekey process. When using the recovery key to unlock the disks, which user has entered his or her recovery key. For information about the number of shares defined and the minimum number required for recovery, issue the encrypt_recovery_key_list command.

Example:

```
encrypt_recovery_key_status
```

Output:

```
Mon Aug 12 20:04:43 IDT 2013
Date Created          User        Status
2013-01-03 18:54:46   secadmin1   Verified
2013-01-03 18:54:46   secadmin2   Verified
2013-01-03 18:54:46   secadmin3   Verified
```

```
2013-01-03 18:54:46   secadmin4   Verified
2013-01-03 19:00:03   secadmin1   Unverified
2013-01-03 19:00:03   secadmin2   Unverified
2013-01-03 19:00:03   secadmin3   Unverified
2013-01-03 19:00:03   secadmin4   Unverified

When entering keys to unlock the disks:
Date Created          User        Status
2013-01-03 19:00:03   secadmin1   Accepted
2013-01-03 19:00:03   secadmin2   Accepted
2013-01-03 19:00:03   secadmin3   Pending
2013-01-03 19:00:03   secadmin4   Pending
```

Field ID	Field output	Default position
create_date	Date Created	1
user	User	2
status	Status	3

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Allowed
Technicians	Allowed

Return codes

- UNSUPPORTED_HARDWARE
  Cannot encrypt on unsupported hardware.

  Troubleshooting: Contact IBM support to verify encryption status.
- CANNOT_READ_FROM_KEY_REPOSITORY
  Failed reading keys from the key repository.

  Troubleshooting: Contact IBM Support.

Recovering key verification

Use the **encrypt_recovery_key_verify** command to confirm that the current user has correctly copied the recovery key share retrieved by the **encrypt_recovery_key_get** command.

```
encrypt_recovery_key_verify key=Key
```

Parameters

Name	Description	Mandatory
key	The 64 character hexadecimal recovery key.	Y

This command is used by security administrators to confirm that they have correctly copied the recovery key presented by the encrypt_recovery_key_get command. Encryption can be enabled (or a rekey can be completed) only when all security administrators have confirmed their respective recovery keys using this command.

Example:

```
encrypt_recovery_key_verify key=B07C4374AC26C4DD3EC2E755EB3FAAF04EC792C8BE0D0CB1C1BAC79998EBEC6D
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Disallowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- UNSUPPORTED_HARDWARE
  Cannot encrypt on unsupported hardware.

  Troubleshooting: Contact IBM support to verify encryption status.
- INVALID_RECOVERY_KEY_FRAGMENT
  The provided recovery key fragment does not match the stored key.

  Troubleshooting: Make sure that the proper key (share) has been used.
- GENERIC_FAILED
  Generic encryption failure.

  Troubleshooting: Contact IBM Support.
- INVALID_RECOVERY_KEY_USER
  The user is not a valid recovery key administrator.

  Troubleshooting: Make sure that the provided usernames are valid.
- CANNOT_READ_FROM_KEY_REPOSITORY
  Failed reading keys from the key repository.

  Troubleshooting: Contact IBM Support.
- RK_FAILED_VERIFY_SLEEP
  Too many failed verify attempts.

  Troubleshooting: Wait a little and try again.
- ENCRYPTION_KR_WRITE_FAILED
  Error writing to the key repository.

  Troubleshooting: Contact IBM Support.
- INVALID_RECOVERY_KEY_STATE
  The recovery key state is inconsistent with the specified option.

  Troubleshooting: Check the recovery key state using the encrypt_recovery_key_status command.
- RECOVERY_KEY_ALREADY_VERIFIED
  The recovery key has already been verified.

  Troubleshooting: Check the recovery key state using the encrypt_recovery_key_list command.

Recovering key share information

Use the **encrypt_recovery_key_list** command to list recovery key share information.

```
encrypt_recovery_key_list
```

This command lists information regarding recovery keys, specifically: How many parts was the recovery key shared across, and how many are needed for the recovery process. When the currently valid recovery keys were created. To retrieve per-user information about the status of each key share, use the encrypt_recovery_key_status command.

Example:

```
encrypt_recovery_key_list
```

Output:

```
Recovery Key Initial Generation:

Date created        Number of Shares    Min Required
------------        ----------------    ------------
2013-03-11 16:00    3                   2

Recovery Key Rekeyed:
Date created        Number of Shares    Min Required
------------        ----------------    ------------
2013-03-11 16:00    3                   2
2013-03-20 16:05    4                   2
```

Field ID	Field output	Default position
create_date	Key Created	1
number_of_shares	Number of Shares	2
min_req	Min Required	3

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Allowed
Technicians	Allowed

Return codes

- UNSUPPORTED_HARDWARE
  Cannot encrypt on unsupported hardware.

  Troubleshooting: Contact IBM support to verify encryption status.
- CANNOT_READ_FROM_KEY_REPOSITORY
  Failed reading keys from the key repository.

  Troubleshooting: Contact IBM Support.

Finishing the recovery process

Use the **encrypt_recovery_finish** command to finish the recovery process and move the system to the On state.

```
encrypt_recovery_finish
```

Upon entering the recovery keys (see Entering a recovery key), this command finishes the recovery process and moves the system to the On state, provided that no more issues exist.

Example:

```
encrypt_recovery_finish
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- ENCRYPT_NOT_ENABLED
  Encryption is not enabled.

  Troubleshooting: Make sure that encryption is enabled and re-run the command.
- RK_ENTER_SYSTEM_STATE_INVALID
  The command is supported in maintenance mode only.

  Troubleshooting: Switch the system state to maintenance mode.

Obtaining a new master key

Use the **encrypt_local_rekey** command to initiate rekeying (getting new cryptographic material) from the local key management.

```
encrypt_local_rekey
```

Example:

```
encrypt_local_rekey
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Disallowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_REKEY
  Are you sure you want to change the local key?

  Troubleshooting: To proceed with issuing the command, enter -y.

Return codes

- ENCRYPT_NOT_ENABLED
  Encryption is not enabled.

  Troubleshooting: Make sure that encryption is enabled and re-run the command.

- CANNOT_GET_XIV_MASTER_KEY
  Problem obtaining XIV master key from the key server.

  Troubleshooting: Make sure that the key service is active and serving keys (obtained from the key server or stored locally). Otherwise, contact IBM Support.

- CANNOT_GET_NEW_KEY_REQUEST
  Error requesting encryption keys from the key server gateway node.

  Troubleshooting: Make sure that the key server is actively serving keys.

- CANNOT_UPDATE_KEY_METADATA
  Cannot update metadata in the key repository for a new key.

  Troubleshooting: Contact IBM Support.

- CANNOT_CANNOT_GENERATE_EXMK_ESKH
  Problem generating EXMK and ESKH.

  Troubleshooting: Contact IBM Support.

- CANNOT_WRITE_TO_KEY_REPOSITORY
  Failed writing keys to the key repository.

  Troubleshooting: Contact IBM Support.

- CANNOT_COPY_KEYS_IN_KEY_REPOSITORY
  Problem copying current keys to the location of the old keys in the key repository.

  Troubleshooting: Contact IBM Support

- ENCRYPTION_KR_WRITE_FAILED
  Error writing to the key repository.

  Troubleshooting: Contact IBM Support.

- ENCRYPTION_KR_READ_FAILED
  Error reading the key repository.

  Troubleshooting: Contact IBM Support.

- NO_LIVE_KEYSERVER_GATEWAY_NODE
  There is no live key server gateway node in the system.

  Troubleshooting: Restart the key server gateway node and try again.

- ENCRYPTION_NOT_IN_LOCAL_SCHEME
  Encryption key management is not set to a local scheme.

  Troubleshooting: Check the output of the encrypt_key_scheme_get command.

Changing the key management scheme

Use the **encrypt_change_key_scheme** command to change the key management scheme.

```
encrypt_change_key_scheme key_scheme=KeyScheme
```

Parameters

Name	Type	Description	Mandatory
key_scheme	Enumeration	Defines which key management scheme (external or local) to use for encryption activation.	Y

This command only supports the change from an external to a local scheme.

Example:

```
encrypt_change_key_scheme key_scheme=local
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Disallowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_CHANGE_KEY_SCHEME
  Are you sure you want to change the key management scheme? This operation is irreversible.

  Troubleshooting: To proceed with issuing the command, enter -y.

Return codes

- ENCRYPT_NOT_ENABLED
  Encryption is not enabled.

  Troubleshooting: Make sure that encryption is enabled and re-run the command.

- CANNOT_GET_XIV_MASTER_KEY
  Problem obtaining XIV master key from the key server.

  Troubleshooting: Make sure that the key service is active and serving keys (obtained from the key server or stored locally). Otherwise, contact IBM Support.

- CANNOT_GET_NEW_KEY_REQUEST
  Error requesting encryption keys from the key server gateway node.

  Troubleshooting: Make sure that the key server is actively serving keys.

- CANNOT_UPDATE_KEY_METADATA
  Cannot update metadata in the key repository for a new key.

  Troubleshooting: Contact IBM Support.

- CANNOT_CANNOT_GENERATE_EXMK_ESKH
  Problem generating EXMK and ESKH.

  Troubleshooting: Contact IBM Support.

- CANNOT_WRITE_TO_KEY_REPOSITORY
  Failed writing keys to the key repository.

  Troubleshooting: Contact IBM Support.

- CANNOT_COPY_KEYS_IN_KEY_REPOSITORY
  Problem copying current keys to the location of the old keys in the key repository.

  Troubleshooting: Contact IBM Support

- ENCRYPTION_KR_WRITE_FAILED
  Error writing to the key repository.

Troubleshooting: Contact IBM Support.

- ENCRYPTION_KR_READ_FAILED
  Error reading the key repository.

  Troubleshooting: Contact IBM Support.

- NO_LIVE_KEYSERVER_GATEWAY_NODE
  There is no live key server gateway node in the system.

  Troubleshooting: Restart the key server gateway node and try again.

- ENCRYPTION_NOT_IN_EXTERNAL_SCHEME
  Encryption key management is not set to the external scheme.

  Troubleshooting: Check the output of the encrypt_key_scheme_get command.

Viewing the key scheme

Use the **encrypt_key_scheme_get** command to view the key scheme defined in the system.

```
encrypt_key_scheme_get
```

Example:

```
encrypt_key_scheme_get
```

Output:

```
Command completed successfully.
encrypt_key_scheme = "LOCAL"
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Allowed
Technicians	Allowed

Multi-site high availability and disaster recovery commands

This section describes the command-line interface (CLI) for Multi-site high availability and disaster recovery.

Important: The input and output syntax of CLI commands uses the legacy terminology of "Master", "SMaster", and "Slave" volumes, which in any documentation except the CLI reference, are referred to as "Primary", "Secondary", and "Tertiary". This inconsistency is a necessary compromise, required to avoid changes to older CLI commands that are in customer use, and also to keep the CLI terminology consistent across the board. The new terminology helps emphasize the commonality between the more recent functions of Multi-site HA/DR, high availability (HyperSwap), and the disaster recovery (Synchronous and Asynchronous mirroring) ones. It is used outside the CLI reference, where broader concepts can be explained.

- **Viewing the Multi-site relation status**
  Use the **multisite_list** command to list the configuration and status of the defined Multi-site relations.
- **Creating a Multi-site relation**
  Use the **multisite_define** command to create a new Multi-site relation.
- **Activating an asynchronous mirror within a Multi-site relation**
  Use the **multisite_activate_async_mirror** command to activate an asynchronous mirror as part of a Multi-site relation.
- **Deleting a Multi-site relation**
  Use the **multisite_delete** command to delete an existing Multi-site relation.
- **Changing roles in a Multi-site relation**
  Use the **multisite_change_role** command to change the role of the local Multi-site relation peer.
- **Adding a Standby mirror to a Multi-site relation**
  Use the **multisite_register_standby_mirror** command to register a Standby mirror in an existing Multi-site relation.
- **Switching roles between Master and SMaster**
  Use the **multisite_switch_roles** command to switch roles between the Master and SMaster volumes or consistency groups.

Viewing the Multi-site relation status

Use the **multisite_list** command to list the configuration and status of the defined Multi-site relations.

```
multisite_list [ vol=VolName | cg=cgName | scope=<cg|volume> ] [ domain=DomainName ]
```

Parameters

Name	Type	Description	Mandatory	Default
scope	Enumeration	All Multi-site relations, volume level Multi-site relations, CG level Multi-site relations.	N	All (if no value is specified)
vol	Object name	Local volume name.	N	[none]
cg	Object name	Local consistency group name.	N	[none]
domain	Object name	The domain name.	N	All Domains.

If no parameters are specified, all the existing Multi-site relations are listed. When a name is indicated, only the specified Multi-site relation is listed.

The following parameters are listed by default:

- Name
- Multisite Object (Volume or CG)
- Multisite ID
- Role (Master, SMaster, Slave)
- State: This parameter can be retrieved only on the Master. The existing options are explained in the table below:

Value	Meaning
`Inactive`	Both two-way relations that make part of a Multi-site relation are inactive.
`Initializing`	One or both two-way relations that make part of a Multi-site relation are in the init state.
`Operational`	The HyperSwap relation is synchronized, automatic failover is active, the asynchronous mirror is RPO OK, the SMaster is ready to assume the ownership of the disaster recovery (DR).
`Degraded`	The asynchronous mirror is RPO lagging.
`Compromised`	This state can be triggered by one of the following: ◦ The connection is disrupted between Master and SMaster and/or Master and Slave, or between SMaster and Slave. The latter is applicable only if a Standby mirror is configured. ◦ Resynchronization between Master and SMaster and/or Master and Slave is in progress. ◦ The machine holding the SMaster cannot assume the ownership of disaster recovery (DR), because the SMaster does not have an updated MRS (most recent snapshot) or LRS (last replicated snapshot.) ◦ A partial role change, either automatic or manual, has occurred between Master and SMaster or Master and Slave. ◦ The Slave cannot perform automatic failover. This may occur because: the Slave is not synchronized, the Slave is disconnected from the Quorum Witness, automatic failover is disabled by the user, or connectivity between Master and Slave has failed.
`Inconsistent`	A relation between Master and Slave is not defined. This can happen when a Standby relation was not created and the SMaster becomes the Master.

- Multisite Standby: The state of the registered Standby mirror between SMaster and Slave. An asynchronous mirroring relation is considered a registered Standby mirror if it is defined on a Multi-site SMaster volume or consistency group, and is inactive and had either been registered or active in the past. The existing options are explained in the table below:

Value	Meaning
`Up`	The Standby mirror is defined, registered and connected.
`Down`	The Standby mirror is defined and registered, but disconnected.
`N/A`	The Standby mirror is either undefined or unregistered.

- Master: The name of the system that holds the Master volume or consistency group. For the local system, local is returned.
- SMaster: The name of the system that holds the SMaster volume or consistency group. For the local system, local is returned.
- Slave: The name of the system that holds the Slave volume or consistency group. For the local system, local is returned.
- Designation: The designated role of the local volume or consistency group in the Multi-site relation.

The following optional parameters can be listed by explicitly specifying the proper columns:

- Main relation active: Indicates whether the two-way relation between the Master and SMaster is active. The available options are Yes, No, or N/A. The latter will appear if the information is returned by the Slave.
- Main relation sync state: Indicates whether the Master and SMaster are synchronized. N/A will appear if the information is returned by the Slave.
- Automatic failover: Indicates whether automatic failover between Master and SMaster is available. Can be retrieved from either Master or SMaster. N/A will appear if the information is returned by the Slave.
- Async mirror active: Indicates whether the asynchronous mirroring between the Master and Slave is active. N/A will appear if the information is returned by the SMaster.
- HA I/O Service: Indicates the local peer's ability to serve I/O. The existing options are Active or Unavailable.
- Standby snapshot state: Indicates whether the SMaster has the data required for the Standby mirror. N/A will appear if the information is returned by the Slave.

Example:

```
multisite_list
```

Output:

```
Name   Multisite Object   Multisite ID       Role     State         Multisite Standby   Master          SMaster   Slave
Designation
------ ------------------ ------------------ -------- ------------- -----------------   --------------- --------- -------
-------- -------------
Va_1   Volume             5BC3473000000000   Master   Operational   Up                  Local           Sync      DR
Master
Vb_2   Volume             5BC3473000000001   SMaster  N/A           Up                  Sync            Local     DR
SMaster
Vc_2   Volume             5BC3473000000002   Slave    N/A           Up                  DR              DR        Local
```

```
Slave
cg_a    CG                  5BC3473000000003    Master    Operational    Up                  Local           Sync      DR
Master
```

Field ID	Field output	Default position
local_peer_name	Name	1
multisite_object	Multisite Object	2
multisite_global_id	Multisite ID	3
multisite_role	Role	4
multisite_state	State	5
standby_snapshot_state	SMaster Mirror Ready	N/A
standby_multisite_state	Multisite Standby	6
master_target_name	Master	7
smaster_target_name	SMaster	8
slave_target_name	Slave	9
designation	Designation	10
master_peer_name	Master Peer	N/A
smaster_peer_name	SMaster Peer	N/A
slave_peer_name	Slave Peer	N/A
main_relation_active	Main relation active	N/A
main_relation_sync_state	Main relation sync state	N/A
automatic_failover	Automatic Failover	N/A
async_mirror_active	Async mirror active	N/A
async_mirror_sync_state	Async mirror sync state	N/A
io_service	HA I/O Service	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Creating a Multi-site relation

Use the **multisite_define** command to create a new Multi-site relation.

```
multisite_define < vol=VolName | cg=cgName >
```

Parameters

Name	Type	Description	Mandatory
vol	Object name	A volume name.	N
cg	Object name	Master consistency group name.	N

Provided that a HyperSwap relation is established between Master and SMaster, and an asynchronous mirror is defined between Master and Slave, issuing this command on the Master creates a named Multi-site relation for a volume or consistency group. If the command completes successfully, the Multi-site relation will be recognized on all the systems involved.

The name of the Multi-site relation on all peers will be identical with the name of the target volume or consistency group on the Master. If a Multi-site relation with this name already exists on any system involved, the command will fail. The command will also fail if any of the following preliminary requirements is not met:

- All three systems are up and connected with the target.
- The target connectivity is configured as follows: Master-SMaster, SMaster-Master, Master-Slave. If the Standby relation is included, then SMaster-Slave connectivity must be configured and up.
- A HyperSwap relation is defined between Master and SMaster and vice versa. The HyperSwap relation status is either Active and Sync, or Inactive.
- An asynchronous mirroring relation is defined between Master and Slave and vice versa. The asynchronous mirroring relation status is either Active and Sync, or Inactive.
- Only one of the relations - either HyperSwap or asynchronous mirroring - is active.
- Either the HyperSwap or the asynchronous mirroring relation is defined as part of Multi-site.

Issuing **multisite_define** on a consistency group, automatically adds the Multi-site relation to each member volume.

Access control

User Category	Permission

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- DESIGNATION_MISMATCH_ARE_YOU_SURE_YOU_WANT_TO_OVERWRITE_EXISTING_DESIGNATION
  There is a mismatch between the designations in a HyperSwap or mirroring relation and the Multi-site roles. Overwrite the existing designations?

Return codes

- TARGET_BAD_NAME
  The target name does not exist.
- VOLUME_BAD_NAME
  The volume name does not exist.
- CONS_GROUP_BAD_NAME
  The consistency group name does not exist.
- MULTISITE_BAD_GLOBAL_ID
  The Multi-site global ID does not exist.
- MULTISITE_SMASTER_TARGET_NOT_CONNECTED
  The Multi-site SMaster target is not connected. The operation cannot be completed.
- MULTISITE_SLAVE_TARGET_NOT_CONNECTED
  The Multi-site Slave target is not connected. The operation cannot be completed.
- MULTISITE_SLAVE_MASTER_RELATION_IS_MISSING
  The relation exists on the Master but not on the Slave.
- MULTISITE_SMASTER_INVALID_CONFIGURATION
  The Multi-site SMaster-Master relation configuration is invalid.
- MULTISITE_ALREADY_DEFINED
  A Multi-site is already defined for this volume/consistency group.
- MULTISITE_MASTER_INVALID_CONFIGURATION
  The Multi-site Master relation configuration is invalid.
- MULTISITE_DEFINE_FAILED_TO_ROLLBACK_MANUAL_CLEANUP_REQUIRED
  An error occurred while defining a Multi-site relation, and the system failed to roll back. Please manually clean up Multi-site relations on all systems.
- MULTISITE_MASTER_INVALID_RELATION_STATE
  An invalid relation state in the Master.

  Troubleshooting: Make sure the existing relations are either both inactive, or one is active/syncronized and the other is inactive.
- MULTISITE_INIT_RELATION_BAD_UID
  The relation exists on the Master, but not on the SMaster or Slave.
- MULTISITE_SMASTER_MASTER_RELATION_IS_MISSING
  The relation exists on the Master, but not on the SMaster.
- MULTISITE_SMASTER_INVALID_RELATION_STATE
  An invalid relation state in the SMaster.

  Troubleshooting: The SMaster-Slave asynchronous mirror must be a Standby relation.
- MULTISITE_INVALID_MASTER_SLAVE_TARGET_CONNECTIVITY
  Cannot define a Multi-site relation because Master-Slave target connectivity at the Master system is invalid. Try again in a few seconds.
- MULTISITE_SLAVE_SMASTER_TARGET_MISMATCH
  The defined SMaster target does not match the defined Slave-SMaster relation target.
- MULTISITE_SLAVE_INVALID_CONFIGURATION
  The Multi-site Slave relation configuration is invalid.
- MULTISITE_SLAVE_INVALID_RELATION_STATE
  An invalid relation state in the Slave.

  Troubleshooting: The SMaster-Slave asynchronous mirror must be a Standby relation.
- TARGET_BAD_TYPE

The target machine is not an XIV machine.

- MULTISITE_SMASTER_MASTER_TARGET_MISMATCH
  The defined Master target does not match the defined SMaster-Master relation target.
- MULTISITE_SLAVE_MASTER_TARGET_MISMATCH
  The defined Master target does not match the defined Slave-Master relation target.
- MULTISITE_INVALID_MASTER_SMASTER_TARGET_CONNECTIVITY
  Cannot define a Multi-site relation because Master-SMaster target connectivity at the Master system is invalid.
- MULTISITE_SMASTER_SLAVE_TARGET_MISMATCH
  The defined Slave target does not match the defined SMaster-Slave relation target.
- VOLUME_HAS_STANDBY_SNAPSHOTS
  A Multi-site volume has standby snapshots created by a previous process.
- MULTISITE_INCOMPATIBLE_TARGET_VERSION
  Multi-site is not supported between the system versions of the specified peers.
- MULTISITE_NUM_OF_MULTISITES_ON_SLAVE_LIMIT_REACHED
  The number of Multi-site relations defined on the Slave system exceeded the limit.
- RELATION_RETRY_OPERATION
  There is an operation in progress on this relation , please try again your request in a few seconds.

  Troubleshooting: Please try again the command in a few seconds.
- MULTISITE_SMASTER_INCOMPATIBLE_VERSION
  The software version on the SMaster system does not support Multi-site.
- MULTISITE_NUM_OF_MULTISITES_ON_SMASTER_LIMIT_REACHED
  The number of Multi-site relations defined on the SMaster system exceeded the limit.
- MULTISITE_SLAVE_INCOMPATIBLE_VERSION
  The software version on the Slave system does not support Multi-site.
- MULTISITE_NUM_OF_MULTISITES_ON_MASTER_LIMIT_REACHED
  The number of Multi-site relations defined on the Master system exceeded the limit.
- MULTISITE_INVALID_MASTER_LOOPBACK_CONFIGURATION
  Cannot define a Multi-site relation with a loopback. A loopback target detected on the Master system.
- MULTISITE_INVALID_SMASTER_LOOPBACK_CONFIGURATION
  Cannot define a Multi-site relation with a loopback. A loopback target detected on the SMaster system.
- MULTISITE_INVALID_SLAVE_LOOPBACK_CONFIGURATION
  Cannot define a Multi-site relation with a loopback. A loopback target detected on the Slave system.
- MULTISITE_MASTER_SMASTER_CLOCK_SKEW_TOO_BIG
  Command failed because the time difference between the Master and SMaster systems is too big.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support
- TARGET_NO_ACCESS
  No access permissions to the secondary machine.
- CONS_GROUP_MEMBER_VOL_IS_MISSING_A_RELATION
  A volume in the consistency group is missing the volume-level Multi-site/mirror/HyperSwap relation.
- MULTISITE_CANNOT_CONTAIN_SYNC_MIRROR
  A Multi-site relation cannot include a synchronous mirroring relation.

Activating an asynchronous mirror within a Multi-site relation

Use the **multisite_activate_async_mirror** command to activate an asynchronous mirror as part of a Multi-site relation.

```
multisite_activate_async_mirror  < vol=VolName | cg=cgName >
```

Parameters

Name	Type	Description	Mandatory
vol	Object name	A volume name.	N
cg	Object name	Master consistency group name.	N

The command can be issued only on the Master volume or consistency group. Issuing it on a standalone Master, that is, on a peer that originally was a Slave and then was promoted to the Master in a failure recovery procedure, is not allowed.

The **multisite_activate_async_mirror** command first ensures that the asynchronous mirror activation will maintain the consistency of the recovery image on the Slave, and then changes the asynchronous mirror state to Active.

If during the preliminary check the command detects that the last consistent snapshot on the Slave is more recent than on the Master, it emits the **MULTISITE_SLAVE_HAS_MORE_RECENT_DATA** return code and fails to activate the asynchronous mirror. In this case, the user may proceed as follows:

- Delete the Multi-site relation, and overwrite the more recent data on the Slave with the Master data.
- Delete the Multi-site relation, and restore the Master with the Slave data.
- Switch roles between Master and SMaster, in case the SMaster has more recent data that matches the Slave data.

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.
- CONS_GROUP_BAD_NAME
  The consistency group name does not exist.
- CONS_GROUP_IS_NOT_MULTISITE
  The local consistency group does not have Multi-site definitions.
- VOLUME_IS_NOT_MULTISITE
  This volume is not part of Multi-site relation.
- SYNC_ALREADY_ACTIVE
  Synchronization is already active.
- MULTISITE_ROLE_IS_NOT_MASTER
  The local peer is not the Master in the Multi-site relation.
- MIRROR_CONFIGURATION_ERROR
  The mirror's local configuration does not match its remote configuration.
- MULTISITE_CONS_GROUP_MEMBERSHIP_MISMATCH
  A Multi-site consistency group contains different volumes on Master and Slave. This problem may occur because a previously issued cg_add_vol or cg_remove_vol command completed only partially.

  Troubleshooting: Please retry the cg_add_vol or cg_remove_vol to ensure the membership in the Multi-site consistency groups in all 3 peers are equivalent
- MIRROR_RETRY_OPERATION
  There is an operation in progress on this mirror.

  Troubleshooting: Retry the command in a few seconds.
- MIRROR_SIZE_MISMATCH
  The secondary and primary volume sizes are different.
- REMOTE_MAX_VOLUMES_REACHED
  The maximum number of volumes on the remote machine is already reached.
- MAX_VOLUMES_REACHED
  The maximum allowed number of volumes is already reached.
- DOMAIN_MAX_VOLUMES_REACHED
  The domain exceeds the maximum allowed number of volumes.
- REMOTE_VOLUME_IS_MASTER
  A volume on the remote machine is already defined as primary.
- MULTISITE_SLAVE_HAS_MORE_RECENT_DATA
  Cannot activate the Multi-site relation because the Slave has more recent data.
- REMOTE_MIRROR_IS_STANDBY
  The remote mirror is marked as Standby.
- MULTISITE_SLAVE_INVALID_STANDBY_CONFIGURATION

The Multi-site Slave standby mirror configuration is invalid.

- MULTISITE_MASTER_INVALID_CONFIGURATION
  The Multi-site Master relation configuration is invalid.
- REMOTE_DOMAIN_MAX_VOLUMES_REACHED
  The maximum number of volumes in the remote machine domain is already reached.
- MULTISITE_ASYNC_MIRROR_IS_NOT_CONNECTED
  Cannot activate a Multi-site relation after role change, if the async mirror is not connected.
- MULTISITE_MASTER_SLAVE_INCONSISTENT_LRS_CONFIGURATION
  Cannot activate the Multi-site relation because the last replicated snapshot (LRS) on the Master is more recent than on the Slave. The snapshot configuration is inconsistent.
- MULTISITE_MASTER_SMASTER_CLOCK_SKEW_TOO_BIG
  Command failed because the time difference between the Master and SMaster systems is too big.
- MULTISITE_ROLE_IS_STANDALONE_MASTER
  Command failed because a Multi-site STANDALONE master cannot be activated.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support
- REMOTE_DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier of the remote system is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support
- REMOTE_SYSTEM_OUT_OF_PHYSICAL_SPACE
  The operation not allowed while the remote system is out of physical space.
- VOLUME_BELONGS_TO_MULTISITE_CONS_GROUP
  A Multi-site volume is part of a Multi-site consistency group.
- VOLUME_BELONGS_TO_MIRRORED_CONS_GROUP
  The volume mirror is part of a consistency group mirror.
- MIRROR_CAN_NOT_BE_ACTIVATED
  Mirroring cannot be activated.
- CONS_GROUP_MEMBER_VOL_IS_MISSING_A_RELATION
  A volume in the consistency group is missing the volume-level Multi-site/mirror/HyperSwap relation.
- REMOTE_MULTISITE_IS_NOT_SMASTER
  The remote peer is not the SMaster.
- MULTISITE_SMASTER_INVALID_CONFIGURATION
  The Multi-site SMaster-Master relation configuration is invalid.
- MULTISITE_BAD_GLOBAL_ID
  The Multi-site global ID does not exist.

Deleting a Multi-site relation

Use the **multisite_delete** command to delete an existing Multi-site relation.

```
multisite_delete < vol=VolName | cg=cgName > [ force=<yes|no> ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	A volume name.	N	N/A
cg	Object name	Master consistency group name.	N	N/A
force	Boolean	Forces the deletion of a Multi-site relation on the local system.	N	no

This command deletes only the Multi-site relation, the two-way relations that make part of it remain intact.

In a healthy configuration, the **multisite_delete** command is issued on the Master volume or consistency group, and affects all three peers. When issued on a consistency group, it automatically deletes the Multi-site relation from each of the member volumes.

If a Multi-site relation cannot be deleted on every peer due a system failure or connectivity problem, it will be deleted partially. This allows deleting a Multi-site configuration when the connectivity between the systems is down.

If a permanent failure has occurred, and the Multi-site configuration cannot be recovered, **multisite_delete** is used to remove the Multi-site relation from the local site, even though the peers may be disconnected. The force=yes parameter can is used in this scenario only.

To delete a Multi-site relation, make sure that only one two-way relation from those defined on the target peer is currently active, either HyperSwap or asynchronous mirror.

When a Multi-site relation is deleted, the Multi-site internal snapshots on the SMaster become external and available for the user to manage. If a new Multi-site relation is created for the same volume or consistency group, those snapshots must be deleted.

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.
- CONS_GROUP_BAD_NAME
  The consistency group name does not exist.
- CONS_GROUP_IS_NOT_MULTISITE
  The local consistency group does not have Multi-site definitions.
- VOLUME_IS_NOT_MULTISITE
  This volume is not part of Multi-site relation.
- MULTISITE_BAD_GLOBAL_ID
  The Multi-site global ID does not exist.
- MULTISITE_TOO_MANY_ACTIVE_RELATIONS
  Only one mirror/HyperSwap relation can be active when a Multi-site relation is being deleted.
- MULTISITE_ROLE_IS_NOT_MASTER
  The local peer is not the Master in the Multi-site relation.
- MULTISITE_DELETE_PARTIAL_FAILURE_MANUAL_CLEANUP_REQUIRED
  Failed to delete the Multi-site relation on all systems. Manual cleanup is required.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support
- VOLUME_BELONGS_TO_MULTISITE_CONS_GROUP
  A Multi-site volume is part of a Multi-site consistency group.
- MIRROR_RETRY_OPERATION
  There is an operation in progress on this mirror.

  Troubleshooting: Retry the command in a few seconds.

Changing roles in a Multi-site relation

Use the **multisite_change_role** command to change the role of the local Multi-site relation peer.

```
multisite_change_role  < vol=VolName | cg=cgName > new_role=<Master|None|Slave|SMaster>
```

Parameters

Name	Type	Description	Mandatory
vol	Object name	A volume name.	N
cg	Object name	Master consistency group name.	N
new_role	Enumeration	Role name of the local Multi-site relation peer.	Y

This command is used in the following cases:

- As part of the procedure restoring the original configuration after an automatic or manual failover
- In order to perform a manual failover when the automatic failover did not take place
- When the recovery from a DR (disaster recovery) site is needed.

multisite_change_role can be issued only for non-operational relations. It is assumed that the command will be issued on all the relevant peers before the relation becomes operational again, so that either the original roles will be restored, or the SMaster and Master will switch roles.

The following role changes are allowed:

- Master to SMaster. For a successful role change from Master to SMaster, the volume can be in any phase, except Initializing. The Master ceases serving host requests, and is set to act as an SMaster and accept replication from the other peer.
- SMaster to Master. For information about the SMaster-to-Master role change, see below.
- Slave to Master. This role change is available only for a standalone Slave, that is not connected to either Master or SMaster. After such Slave becomes Master, it cannot be incorporated into an existing Multi-site relation. Instead, a new Multi-site relation must be defined.

The SMaster-to-Master role change

Preliminary requirements

- Make sure that the original Master is not available and cannot become available while the other peer is acting as the Master. To verify this, run the **multisite_list** command on the original Master and check the value of the HA I/O Service output field. Only if the returned value is Unavailable, proceed with issuing the **multisite_change_role** command on the SMaster.
- Stop the applications that use Multi-site volume(s) or consistency group(s). Note that the applications must be stopped completely and not merely paused. This is to make sure that an application does not use any cached state when accessing the volume or consistency group for the first time after the role change.

Important: A failure to comply with either of the above requirements may result in a data integrity issue.

Data consistency on the SMaster

After an SMaster is successfully changed to a Master, the volume (or consistency group) starts accepting requests from hosts. Upon explicit activation, it starts replicating to the other peer (the original Master).

If HyperSwap is interrupted in the middle of the re-synchronization process, the SMaster may be inconsistent. The last consistent image of the SMaster volume (or consistency group) is preserved in the last consistent snapshot (LCS), which is automatically created before the re-synchronization starts. If the LCS exists, the command emits a **ARE_YOU_SURE_YOU_WANT_TO_CHANGE_MULTISITE_ROLE_WITH_LCS** warning. In this case, before issuing the before issuing the **multisite_change_role** command, the administrator must choose whether to use the existing contents of the original SMaster, which may be inconsistent, or revert the original SMaster to the LCS.

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- SOME_DATA_WILL_BE_LOST_ARE_YOU_SURE
  Are you sure you want the mirror/HyperSwap local peer to become secondary and lose the data that was not replicated?
- ARE_YOU_SURE_YOU_WANT_TO_CHANGE_MULTISITE_ROLE
  Are you sure you want to change the Multi-site role?
- ARE_YOU_SURE_YOU_WANT_TO_CHANGE_MULTISITE_ROLE_WITH_LCS
  Are you sure you want to change the Multi-site role with the existing last consistent snapshot?

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.
- CONS_GROUP_BAD_NAME
  The consistency group name does not exist.
- CONS_GROUP_IS_NOT_MULTISITE
  The local consistency group does not have Multi-site definitions.
- VOLUME_IS_NOT_MULTISITE
  This volume is not part of Multi-site relation.
- MULTISITE_SMASTER_CAN_ONLY_BE_CHANGED_TO_MASTER
  On this system, the Multi-site SMaster can be changed only to the Master.
- MULTISITE_SLAVE_CAN_ONLY_BE_CHANGED_TO_MASTER
  On this system, the Multi-site Slave can be changed only to the Master.
- MULTISITE_MASTER_CAN_ONLY_BE_CHANGED_TO_SLAVE

On this system, the Multi-site Master can be changed only to the Slave.

- MIRROR_RETRY_OPERATION
  There is an operation in progress on this mirror.

  Troubleshooting: Retry the command in a few seconds.

- HA_RETRY_OPERATION
  An operation is in progress on this HyperSwap relation.

  Troubleshooting: Try issuing the command again in a few seconds.

- MULTISITE_MASTER_CAN_ONLY_BE_CHANGED_TO_SMASTER
  On this system, the Multi-site Master can be changed only to the SMaster.
- MIRROR_IS_INITIAL
  The operation is not permitted during the Initialization phase.
- HA_IS_INITIAL
  The operation is not permitted during the HyperSwap relation initialization phase.
- MIRROR_HAS_NO_SYNCHED_SNAPSHOT
  The mirror does not have a synchronized snapshot.
- HA_HAS_NO_SYNCHED_SNAPSHOT
  This HyperSwap volume does not have a synchronized snapshot.
- MASTER_CANNOT_BE_DEMOTED
  The primary volume cannot be demoted to secondary. Peer status mismatch.
- MIRROR_IS_ACTIVE
  Remote mirroring is currently active.
- HA_IS_ACTIVE
  HyperSwap is currently active.
- VOLUME_HAS_DATA_MIGRATION
  Data Migration is defined for this volume.
- COMMAND_NOT_SUPPORTED_FOR_OLVM_VOLUMES
  This command is not supported for IBM Hyper-Scale Mobility volumes.
- VOLUME_BELONGS_TO_MULTISITE_CONS_GROUP
  A Multi-site volume is part of a Multi-site consistency group.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support

Adding a Standby mirror to a Multi-site relation

Use the **multisite_register_standby_mirror** command to register a Standby mirror in an existing Multi-site relation.

```
multisite_register_standby_mirror < vol=VolName | cg=cgName >
```

Parameters

Name	Type	Description	Mandatory
vol	Object name	A volume name.	N
cg	Object name	Master consistency group name.	N

Issuing this command is only allowed on the system that holds the SMaster.

As a preliminary requirement for **multisite_register_standby_mirror**, the user must create the SMaster-Slave asynchronous mirroring relation.

multisite_register_standby_mirror verifies the configuration of the existing SMaster-Slave mirroring relation and adds it as a registered Standby to the Multi-site relation.

When this command is applied to a consistency group, the Standby relation is registered on every volume in the consistency group.

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed

User Category	Permission
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.
- CONS_GROUP_BAD_NAME
  The consistency group name does not exist.
- CONS_GROUP_IS_NOT_MULTISITE
  The local consistency group does not have Multi-site definitions.
- VOLUME_IS_NOT_MULTISITE
  This volume is not part of Multi-site relation.
- TARGET_BAD_NAME
  The target name does not exist.
- MULTISITE_ROLE_IS_NOT_SMASTER
  The local peer is not the SMaster in the Multi-site relation.
- MULTISITE_SMASTER_INVALID_CONFIGURATION
  The Multi-site SMaster-Master relation configuration is invalid.
- MULTISITE_SLAVE_INVALID_STANDBY_CONFIGURATION
  The Multi-site Slave standby mirror configuration is invalid.
- MULTISITE_STANDBY_RELATION_ALREADY_DEFINED
  A standby relation is already defined on this system.
- MULTISITE_SLAVE_SMASTER_TARGET_MISMATCH
  The defined SMaster target does not match the defined Slave-SMaster relation target.
- MULTISITE_SLAVE_INVALID_CONFIGURATION
  The Multi-site Slave relation configuration is invalid.
- MULTISITE_SLAVE_INVALID_RELATION_STATE
  An invalid relation state in the Slave.

  Troubleshooting: The SMaster-Slave asynchronous mirror must be a Standby relation.
- MULTISITE_SMASTER_INVALID_RELATION_STATE
  An invalid relation state in the SMaster.

  Troubleshooting: The SMaster-Slave asynchronous mirror must be a Standby relation.
- MULTISITE_SMASTER_INVALID_STANDBY_CONFIGURATION
  The standby relation configuration on SMaster is invalid.
- MULTISITE_SMASTER_SLAVE_TARGET_MISMATCH
  The defined Slave target does not match the defined SMaster-Slave relation target.
- MULTISITE_INVALID_SMASTER_LOOPBACK_CONFIGURATION
  Cannot define a Multi-site relation with a loopback. A loopback target detected on the SMaster system.
- MULTISITE_INVALID_SLAVE_LOOPBACK_CONFIGURATION
  Cannot define a Multi-site relation with a loopback. A loopback target detected on the Slave system.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support
- CONS_GROUP_MEMBER_VOL_IS_MISSING_A_RELATION
  A volume in the consistency group is missing the volume-level Multi-site/mirror/HyperSwap relation.
- VOLUME_BELONGS_TO_MULTISITE_CONS_GROUP
  A Multi-site volume is part of a Multi-site consistency group.

Switching roles between Master and SMaster

Use the **multisite_switch_roles** command to switch roles between the Master and SMaster volumes or consistency groups.

```
multisite_switch_roles <vol=VolName | cg=cgName> [ verify_standby=<yes|no> ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	Local volume name.	N	N/A
cg	Object name	Local consistency group name.	N	N/A
verify_standby	Boolean	If set to yes, the command can be successfully completed only if the SMaster-Slave asynchronous Standby mirror can be activated and the SMaster can undertake the disaster recovery. If set to no, the command's successful completion is not contingent on the Standby mirror state.	N	no

This command can only be issued on the Master volume or consistency group. The command can be successfully completed only if:

- the HyperSwap relation is synchronized
- the asynchronous mirroring relation between the Master and the Slave is deactivated or not configured

In addition, if verify_standby is set to yes, the command verifies whether the SMaster-Slave asynchronous Standby mirror can be activated and the SMaster can undertake the disaster recovery.

Important:
Conditions may change after the verification. As a result, the subsequent activation attempt may still fail even though all conditions were verified before switching the roles.

If all the requirements are met, **multisite_switch_roles** allows the system to perform all pending writes. Only after they have been committed, the roles are switched.

Following the execution of the command:

- The volume or consistency group that was previously the Master becomes the SMaster
- The volume or consistency group that was previously the SMaster becomes the Master
- The Multi-site relation remains active
- The asynchronous mirroring relation between the SMaster and Slave is considered registered standby, provided that it had been active and then deactivated in preparation for the role switch

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_SWITCH_ROLES
  Are you sure you want to switch the roles in this relation?
- ARE_YOU_SURE_YOU_WANT_TO_SWITCH_THE_PEER_ROLES
  The system that owns the primary volume is currently not connected to the Quorum Witness. If you switch the peer roles, the automatic failover may become unavailable. Are you sure you want to continue?

Return codes

- VOLUME_BAD_NAME
  The volume name does not exist.
- CONS_GROUP_BAD_NAME
  The consistency group name does not exist.
- MULTISITE_ROLE_IS_NOT_MASTER
  The local peer is not the Master in the Multi-site relation.
- VOLUME_BELONGS_TO_MULTISITE_CONS_GROUP
  A Multi-site volume is part of a Multi-site consistency group.
- MULTISITE_STANDBY_MIRROR_NOT_REGISTERED
  A standby mirror is either not registered or not defined for this Multi-site relation.
- MULTISITE_STANDBY_MIRROR_IS_NOT_CONNECTED
  There is currently no connection between the SMaster and the standby async mirror target system.

  Troubleshooting: The relation is part of a Multi-site relation, therefore connectivity between the SMaster and the standby async mirror Slave is required
- MULTISITE_ASYNC_MIRROR_IS_ACTIVE
  The mirror between the Multi-site Master and Slave is active.

Troubleshooting: The relation in which the roles must be switched is part of a Multi-site relation. The async mirror on the Master must be deactivated.

- MULTISITE_SMASTER_HAS_NO_SYNCHED_SNAPSHOT
  The Multi-site SMaster does not have a synchronized snapshot.
- HA_IS_NOT_SYNCHRONIZED
  The HyperSwap relation is not synchronized.
- VOLUME_HAS_DATA_MIGRATION
  Data Migration is defined for this volume.
- REMOTE_TARGET_NOT_CONNECTED
  There is currently no connection from the target system.
- HA_CONNECTIVITY_NOT_SUFFICIENT
  The connectivity between the systems is not sufficient for the automatic failover.
- HA_RELATION_MASTER_COULD_NOT_UPDATE_QW_AFTER_RETURN_TO_GOOD_STATE
  The primary volume had control over the relationship during a past failure, and was unable to update the Quorum Witness after the recovery.
- HA_HAS_SYNC_JOB
  This operation is not permitted on a HyperSwap relation with active sync jobs.
- HA_RETRY_OPERATION
  An operation is in progress on this HyperSwap relation.

  Troubleshooting: Try issuing the command again in a few seconds.
- REMOTE_HA_IS_NOT_ACTIVE
  The remote peer in this HyperSwap relation is not active.
- COMMAND_NOT_SUPPORTED_FOR_OLVM_VOLUMES
  This command is not supported for IBM Hyper-Scale Mobility volumes.
- DATA_REDUCTION_TIER_IS_OFFLINE
  The data reduced tier is offline, the operation is not allowed.

  Troubleshooting: Contact IBM Support
- SYSTEM_OUT_OF_PHYSICAL_SPACE
  The operation not allowed while the system is out of physical space.
- VOLUME_IS_NOT_MULTISITE
  This volume is not part of Multi-site relation.
- CONS_GROUP_IS_NOT_MULTISITE
  The local consistency group does not have Multi-site definitions.
- VOLUME_BELONGS_TO_HA_CONS_GROUP
  This volume belongs to a HyperSwap consistency group. The operation can be carried out on a volume which is not part of a consistency group, or on the consistency group itself.
- REMOTE_MAY_NOT_HAVE_COMPLETED_THE_OPERATION
  The operation may be not yet completed on the remote target.
- MIRROR_RETRY_OPERATION
  There is an operation in progress on this mirror.

  Troubleshooting: Retry the command in a few seconds.

Security configuration commands

This chapter describes the command line interface (CLI) for security configuration.

- **Listing configuration parameters for a communication protocol**
  Use the **protocol_config_list** command to list configuration parameters per communication protocol.
- **Setting configuration parameters for a communication protocol**
  Use the **protocol_config_set** command to set configuration parameters for a specific communication protocol or for all protocols.

Listing configuration parameters for a communication protocol

Use the **protocol_config_list** command to list configuration parameters per communication protocol.

```
protocol_config_list [ protocol=<xcli|kmip|cim> ]
```

Parameters

Name	Type	Description	Mandatory	Default
protocol	Enumeration	The available options are: XCLI, KMIP, CIM, or All (if no value is specified).	N	All (if no value is specified).

The following default parameters are shown:

- Protocol Name (XCLI, KMIP, CIM)
- Minimal TLS Version (TLS1.0, TLS1.1, TLS1.2)

Example:

```
protocol_config_list
```

Output:

```
Protocol Name            Minimal TLS Version
------------------------ -----------------------------------
CIM                      TLS1.2
KMIP                     TLS1.2
XCLI                     TLS1.2
```

Field ID	Field output	Default position
protocol_name	Protocol Name	1
min_tls_level	Minimal TLS Version	2

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Allowed
Read-only users	Allowed
Technicians	Allowed

Setting configuration parameters for a communication protocol

Use the **protocol_config_set** command to set configuration parameters for a specific communication protocol or for all protocols.

```
protocol_config_set [ protocol=<xcli|kmip|cim> ] min_tls_level=<tls1.0|tls1.1|tls1.2>
```

Parameters

Name	Type	Description	Mandatory	Default
protocol	Enumeration	The available options are: XCLI, KMIP, CIM (case insensitive), or All (if no value is specified).	N	All (if no value is specified).
min_tls_level	Enumeration	Minimum TLS version to support. The available options are: TLS1.0, TLS1.1, or TLS1.2 (case insensitive).	Y	N/A

Example:

```
protocol_config_set protocol=KMIP min_tls_level=TLS1.2
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Events

VOLUME_CREATE

Severity	Description
informational	Volume was created with name '*volume.name*' and size *volume.size*GB in Storage Pool with name '*volume.pool_name*'.

VOLUME_CREATE_FAILED_TOO_MANY

Severity	Description	Troubleshooting
warning	Volume with name '*name*' could not be created. You are attempting to add more volumes than the system permits.	Delete volumes to allow new ones to be created.

VOLUME_RENAME

Severity	Description
informational	Volume with name '*old_name*' was renamed '*volume.name*'.

SUBORDINATE_VOLUME_RENAME

Severity	Description
informational	Remote volume with name '*old_name*' was renamed '*volume.name*' by local system.

VOLUME_RESIZE

Severity	Description
informational	Volume with name '*volume.name*' was resized from *old_size*GB to *volume.size*GB.

SECONDARY_VOLUME_RESIZE

Severity	Description
informational	Secondary volume with name '*volume.name*' was resized by primary machine from *old_size*GB to *volume.size*GB.

VOLUME_DELETE

Severity	Description
informational	Volume with name '*volume.name*' was deleted and its data is no longer accessible.

VOLUME_FORMAT

Severity	Description
informational	Volume with name '*volume.name*' was formatted.

VOLUME_COPY

Severity	Description
informational	Volume with name '*source.name*' was copied to volume with name '*target.name*'.

VOLUME_COPY_DIFF

Severity	Description
informational	Volume with name '*source.name*' was diff-copied from base '*base.name*' to volume with name '*target.name*'.

VOLUME_LOCK

Severity	Description
informational	Volume with name '*volume.name*' was locked and set to 'read-only'.

VOLUME_UNLOCK

Severity	Description
informational	Volume with name '*volume.name*' was unlocked and set to 'writable'.

SUBORDINATE_VOL_LOCK

Severity	Description
informational	Remote volume with name '*volume.name*' was locked and set to 'read-only' by local machine'.

SUBORDINATE_VOL_UNLOCK

Severity	Description
informational	Remote volume with name '*volume.name*' was unlocked and set to 'writable' by local machine.

VOLUME_MOVE

Severity	Description
informational	Volume with name '*volume.name*' has been moved from Storage Pool '*orig_pool.name*' to Pool '*pool.name*'.

OLVM_CREATE

Severity	Description
informational	Volume with name '*volume.name*' was defined as an IBM Hyper-Scale Mobility.

OLVM_OWNER_CREATE

Severity	Description
informational	Volume with name '*volume.name*' was defined as an OWNER IBM Hyper-Scale Mobility.

OLVM_OWNER_ACTIVATED

Severity	Description
informational	IBM Hyper-Scale Mobility Owner Volume with name '*volume.name*' was activated.

OLVM_ACTIVATE

Severity	Description
informational	IBM Hyper-Scale Mobility Volume with name '*volume.name*' was activated.

OLVM_DEACTIVATE

Severity	Description
informational	IBM Hyper-Scale Mobility Volume with name '*volume.name*' was deactivated.

OLVM_REMOTE_ACTIVATE

Severity	Description
informational	IBM Hyper-Scale Mobility Volume with name '*volume.name*' was activated.

OLVM_REMOTE_DEACTIVATE

Severity	Description
informational	IBM Hyper-Scale Mobility Volume with name '*volume.name*' was deactivated.

OLVM_PROXY_INITIATED

Severity	Description
informational	IBM Hyper-Scale Mobility Volume with name '*volume.name*' move to Proxy state started.

OLVM_PROXY

Severity	Description
informational	IBM Hyper-Scale Mobility Volume with name '*volume.name*' entered Proxy state. Volume data on the system is freed.

OLVM_DELETE

Severity	Description	Troubleshooting
informational	IBM Hyper-Scale Mobility Volume process with name '*name*' was deleted.	Delete volumes to allow new ones to be created.

OLVM_ABORT

Severity	Description
informational	IBM Hyper-Scale Mobility Volume process with name '*volume.name*' was aborted.

OLVM_OWNER_DELETE

Severity	Description
informational	IBM Hyper-Scale Mobility Owner Volume process with name '*volume.name*' was deleted.

OLVM_OWNER_ABORT

Severity	Description
informational	IBM Hyper-Scale Mobility Owner Volume process with name '*volume.name*' was aborted.

OLVM_SYNC_STARTED

Severity	Description
informational	IBM Hyper-Scale Mobility Synchronization of volume '*name*' has started.

OLVM_SYNC_ENDED

Severity	Description
informational	IBM Hyper-Scale Mobility Synchronization of volume '*name*' has ended.

QUORUM_WITNESS_DEFINED

Severity	Description
informational	Quorum witness '*Quorum Witness Name*' defined.

QUORUM_WITNESS_DELETED

Severity	Description
informational	Quorum witness '*Quorum Witness Name*' deleted.

QUORUM_WITNESS_UPDATED

Severity	Description
informational	Quorum witness '*Quorum Witness Name*' updated.

QUORUM_WITNESS_RENAMED

Severity	Description
informational	Quorum witness with name '*Old Name*' was renamed '*New Name*'.

QUORUM_WITNESS_ACTIVATION_START

Severity	Description
informational	Quorum witness '*Quorum Witness Name*' activation started.

QUORUM_WITNESS_ACTIVATION_SUCCESSFUL

Severity	Description
informational	Quorum witness '*Quorum Witness Name*' activated successfully.

QUORUM_WITNESS_ACTIVATION_FAILED

Severity	Description
minor	Quorum witness '*Quorum Witness Name*' activation failed.

QUORUM_WITNESS_DEACTIVATION_START

Severity	Description
informational	Quorum witness '*Quorum Witness Name*' deactivation started.

QUORUM_WITNESS_DEACTIVATION_SUCCESSFUL

Severity	Description
informational	Quorum witness '*Quorum Witness Name*' deactivated successfully.

QUORUM_WITNESS_DEACTIVATION_FAILED

Severity	Description
minor	Quorum witness '*Quorum Witness Name*' deactivation failed.

QUORUM_WITNESS_COMMUNICATION_DOWN

Severity	Description
major	Communication with Quorum Witness '*Quorum Witness Name*' is down.

QUORUM_WITNESS_COMMUNICATION_UP

Severity	Description
informational	Communication with Quorum Witness '*Quorum Witness Name*' is up.

QUORUM_WITNESS_HEARTBEATS_OPERATIONAL

Severity	Description
informational	Successfully sending heartbeats to Quorum Witness '*Quorum Witness Name*'.

QUORUM_WITNESS_HEARTBEATS_FAILING

Severity	Description
major	Failing to send heartbeats to Quorum Witness '*Quorum Witness Name*'.

QUORUM_WITNESS_CERTIFICATE_ABOUT_TO_EXPIRE

Severity	Description
variable	Quorum witness '*Quorum Witness Name*' certificate is about to expire on *Expiration Date* (*Counter* notification).

QUORUM_WITNESS_CERTIFICATE_HAS_EXPIRED

Severity	Description
critical	Quorum witness '*Quorum Witness Name*' certificate has expired.

QUORUM_WITNESS_EXTERNAL_NAME_CHANGED

Severity	Description
informational	Quorum witness '*Quorum Witness Name*' has acquired a new external name '*External Name*'.

QUORUM_WITNESS_EVENT_REPORT

Severity	Description
variable	Event reported from Quorum Witness '*Quorum Witness Name*': *Event Description* (*System ID*).

QUORUM_WITNESS_MISSING_EVENTS

Severity	Description
informational	The events from '*First*' through '*Last*' are missing from Quorum Witness '*Quorum Witness Name*'.

QUORUM_WITNESS_DB_RECOVERY_NEEDED

Severity	Description
critical	Quorum witness '*Quorum Witness Name*' db has failed. Recovery is needed.

QUORUM_WITNESS_DB_RESTORE_PENDING

Severity	Description
critical	Quorum witness '*Quorum Witness Name*' db restore is pending. The db will be restored when the Quorum Witness is activated.

QUORUM_WITNESS_DB_RECOVERED

Severity	Description
major	Quorum witness '*Quorum Witness Name*' db was successfully recovered.

QUORUM_WITNESS_LOG_AUTHENTICATION_SET

Severity	Description
informational	Quorum witness '*Quorum Witness Name*' authentication for log retrieval was set.

QUORUM_WITNESS_LOG_AUTHENTICATION_SET_FAILED

Severity	Description
minor	Quorum witness '*Quorum Witness Name*' authentication for log retrieval setup failed, '*Failure Reason*'.

QUORUM_WITNESS_LOG_AUTHENTICATION_RESET

Severity	Description
informational	Quorum witness '*Quorum Witness Name*' authentication for log retrieval was reset.

QUORUM_WITNESS_LOG_AUTHENTICATION_RESET_FAILED

Severity	Description
minor	Quorum witness '*Quorum Witness Name*' authentication reset for log retrieval setup failed, '*Failure Reason*'.

HA_CREATE

Severity	Description
informational	A HA relation was defined for Volume '*local volume name*' on Target '*target name*'.

HA_CREATE_SLAVE

Severity	Description
informational	A HA relation was defined by Target '*target name*' for Volume '*local volume name*'.

HA_ACTIVATE

Severity	Description
informational	The HA relation of peer '*local peer name*' on Target '*target name*' was activated.

HA_SLAVE_ACTIVATE

Severity	Description
informational	The HA relation on peer '*local peer name*' was activated by target '*target name*'.

HA_SLAVE_DEACTIVATE

Severity	Description
informational	The HA relation on peer '*local peer name*' was deactivated by target '*target name*'.

HA_INCOMPATIBLE_VERSION_FOR_UNMAP_SUPPORT

Severity	Description
warning	A HA of peer '*local peer name*' on target '*target name*' cannot support unmap, remote machine has incompatible version.

CG_HA_CREATE

Severity	Description
informational	High availability mirror was defined for Consistency Group '*local CG name*' on Target '*target name*'. Remote Consistency Group is '*remote CG name*'.

CG_HA_CREATE_SLAVE

Severity	Description
informational	High availability mirror was defined by Target '*target name*' for CG '*local CG name*'. Remote CG is '*remote CG name*'.

HA_SYNC_STARTED

Severity	Description

Severity	Description
informational	Synchronization of remote ha relation of volume '*local volume name*' on Target '*target name*' has started.

HA_SYNC_ENDED

Severity	Description
informational	Synchronization of remote ha relation of peer '*local peer name*' on target '*target name*' has ended.

HA_DEACTIVATE

Severity	Description
informational	The HA relation of peer '*local peer name*' on Target '*target name*' was deactivated.

HA_DELETE

Severity	Description
informational	The HA relation of peer '*local peer name*' to a peer on Target '*target name*' was deleted.

HA_SWITCH_ROLES_TO_SLAVE

Severity	Description
informational	Local peer '*local peer name*' is now Slave of a peer on Target '*target name*'.

HA_SWITCH_ROLES_TO_MASTER

Severity	Description
informational	Local peer '*local peer name*' is now Master of a peer on Target '*target name*'.

HA_REVERSE_ROLE_TO_SLAVE

Severity	Description
informational	Local peer '*local peer name*' is now Slave of a peer on Target '*target name*'.

HA_REVERSE_ROLE_TO_MASTER

Severity	Description
informational	Local peer '*local peer name*' is now Master of a peer on Target '*target name*'.

HA_AVAILABILITY_RESTORED

Severity	Description
informational	Local peer '*local peer name*' was made available.

HA_SNAPSHOT_CREATE

Severity	Description
informational	HA Snapshot named '*snapshot.name*' was created for volume named '*volume.name*'.

HA_SNAPSHOT_CREATE_FAILED

Severity	Description
minor	HA Remote snapshot named '*snapshot name*' was not created successfully. Error code is '*error*'

HA_SNAPSHOT_OVERWRITE

Severity	Description
informational	HA Snapshot named '*snapshot.name*' was overwritten for volume named '*volume.name*'.

HA_SLAVE_SNAPSHOT_CREATE

Severity	Description
informational	HA Snapshot named '*snapshot.name*' was created for volume named '*volume.name*'.

HA_SLAVE_SNAPSHOT_OVERWRITE

Severity	Description
informational	HA Snapshot named '*snapshot.name*' was overwritten for volume named '*volume.name*'.

HA_HIGH_AVAILABILITY_ENABLED

Severity	Description
informational	HA relation on peer '*local peer name*' high availability is enabled.

HA_HIGH_AVAILABILITY_DISABLED

Severity	Description
informational	HA relation on peer '*local peer name*' high availability is disabled by user.

HA_CONVERTED_INTO_MIRROR

Severity	Description
informational	HA relation of peer '*local peer name*' on Target '*target name*' was converted into mirror.

HA_CONVERTED_INTO_MIRROR_SLAVE

Severity	Description
informational	HA relation of peer '*local peer name*' on Target '*target name*' was converted into mirror.

TARGET_CONNECTION_HA_SUFFICIENT

Severity	Description
informational	Target named '*target.name*' is ha sufficient connected.

TARGET_CONNECTION_HA_INSUFFICIENT

Severity	Description
critical	Target named '*target.name*' is ha insufficient connected.

TARGET_IS_HA_HEALTHY

Severity	Description
informational	Target named '*target.name*' is HA healthy according to Quorum Witness.

TARGET_IS_HA_UNHEALTHY

Severity	Description
major	Target named '*target.name*' is HA unhealthy according to Quorum Witness.

TARGET_DATA_SERVICE_FAILURE

Severity	Description
major	Target named '*target.name*' has data service failure. Reason: *Reason*.

TARGET_RESUMED_NORMAL_OPERATION

Severity	Description
informational	Target named '*target.name*' resumed normal operation.

HA_AUTOMATIC_FAILOVER_SUCCESSFUL

Severity	Description
warning	HA Slave relation on peer '*local peer name*' has completed failover.

HA_MASTER_REMAINS_AVAILABLE

Severity	Description
informational	HA Master relation on peer '*local peer name*' remains available. Reason: *Reason*.

HA_MASTER_AVAILABLE

Severity	Description
informational	HA Master relation on peer '*local peer name*' is available.

HA_MASTER_UNAVAILABLE

Severity	Description
major	HA Master relation on peer '*local peer name*' is unavailable. Reason: *Unavailable Reason*.

HA_SLAVE_AVAILABLE

Severity	Description
informational	HA Slave relation on peer '*local peer name*' is available.

HA_SLAVE_UNAVAILABLE

Severity	Description
major	HA Slave relation on peer '*local peer name*' is unavailable. Reason: *Unavailable Reason*.

HA_MASTER_RELEASED_CONTROL

Severity	Description
informational	HA Master released control on relation '*local peer name*'.

HA_SLAVE_RELEASED_CONTROL

Severity	Description
informational	HA Slave released control on relation '*local peer name*'.

MIRROR_CONVERT_INTO_HA_STARTED

Severity	Description
informational	Remote mirror of peer '*local peer name*' on Target '*target name*' conversion into HA started.

MIRROR_CONVERT_INTO_HA_STARTED_SLAVE

Severity	Description
informational	Remote mirror of peer '*local peer name*' on Target '*target name*' conversion into HA started.

MIRROR_CONVERT_INTO_HA_ENDED

Severity	Description
informational	Remote mirror of peer '*local peer name*' on Target '*target name*' conversion into HA ended.

MIRROR_CONVERT_INTO_HA_ENDED_SLAVE

Severity	Description
informational	Remote mirror of peer '*local peer name*' on Target '*target name*' conversion into HA ended.

HA_REESTABLISH_FAILED_CONFIGURATION_ERROR

Severity	Description	Troubleshooting
major	HyperSwap reestablish failed. Either configuration of remote HyperSwap of peer '*local peer name*' on target '*target name*' does not match local configuration.	Make sure configuration on both machines is compatible and activate the HyperSwap. If problem persists contact IBM support.

HA_END_SYNC_FAILED_CONFIGURATION_ERROR

Severity	Description	Troubleshooting
major	Configuration of remote HyperSwap of peer '*local peer name*' on target '*target name*' does not match local configuration.	Make sure configuration on both machines is compatible and activate the HyperSwap. If problem persists contact IBM support.

HA_CHANGE_DESIGNATION

Severity	Description
informational	Local peer '*local peer name*' switched its designated role with peer on Target '*target name*'. It is now *designation*.

TARGET_QUORUM_WITNESS_NOT_LAPSED

Severity	Description
informational	Target '*target.name*' Quorum Witness '*quorum_witness.name*' is not lapsed.

TARGET_QUORUM_WITNESS_LAPSED

Severity	Description
Critical	Target '*target.name*' Quorum Witness '*quorum_witness.name*' is lapsed.

HA_VOLUME_MAPPING_MISSING

Severity	Description	Troubleshooting
major	The HyperSwap volume '*local volume name*' must be mapped to allow a host to continue to perform I/O after an automatic failover.	Map the volume to the relevant host.

HA_VOLUME_MAPPING_OK

Severity	Description
informational	The HyperSwap volume '*local volume name*' is now mapped to a host.

DATA_REBUILD_STARTED

Severity	Description
informational	Rebuild process started because system data is not protected. *data_percent*% of the data must be rebuilt.

DATA_REBUILD_COMPLETED

Severity	Description
informational	Rebuild process completed. System data is now protected.

DATA_REDIST_STARTED

Severity	Description
informational	Starting data transfer to new disks.

DATA_REDIST_COMPLETED

Severity	Description
informational	Completed data transfer to new disks.

DATA_REBUILD_COMPLETED_REDIST_STARTED

Severity	Description
informational	Rebuild process completed. System data is now protected. Starting data transfer to new disks.

DATA_REDIST_TIME_LIMIT_EXCEEDED

Severity	Description
warning	Data redistribution is taking too long. *data_percent*% of the required redistribution still remains to be done.

DATA_REDIST_BLOCKED

Severity	Description
informational	Blocking data transfer to new phased in media for *delay_in_seconds* seconds.

DATA_REDIST_ALLOWED

Severity	Description
informational	Allowing data transfer to new phased in media.

STORAGE_POOL_EXHAUSTED

Severity	Description	Troubleshooting
major	Pool '*pool*' is full. All volumes are locked.	Enlarge Storage Pool or move or delete volumes or Clones with Clone Deletion Priority 0.

STORAGE_POOL_UNLOCKED

Severity	Description
major	Pool '*pool*' has empty space. All volumes are unlocked.

STORAGE_POOL_SNAPSHOT_USAGE_INCREASED

Severity	Description
variable	Usage by snapshots of Storage Pool with name '*pool.name*' has reached *current*%.

STORAGE_POOL_SNAPSHOT_USAGE_DECREASED

Severity	Description
informational	Usage by snapshots of Storage Pool with name '*pool.name*' has decreased to *current*%.

HOST_CONNECTED

Severity	Description
informational	Host '*host*' has connected to the system.

HOST_DISCONNECTED

Severity	Description
warning	Host '*host*' has disconnected from the system.

HOST_MULTIPATH_OK

Severity	Description
informational	Host '*host*' has redundant connections to the system. #paths=*npaths*

HOST_NO_MULTIPATH_ONLY_ONE_PORT

Severity	Description
warning	Host '*host*' is connected to the system through only one of its ports. #paths=*npaths*

HOST_MULTIPATH_NOT_RESILIENT_TO_MODULES_FAILURES

Severity	Description	Troubleshooting
informational	Host '*host*' is connected to the system through less than three modules and hence is less resilient to modules failures.	Add host paths to the system through at least three active modules.

HOST_NO_MULTIPATH_ONLY_ONE_MODULE

Severity	Description
informational	Host '*host*' is connected to the system through only one interface module. #paths=*npaths*

SYSTEM_CAPACITY_USAGE_INCREASED

Severity	Description
variable	Usage of system physical capacity increased to *current*%.

SYSTEM_CAPACITY_USAGE_DECREASED

Severity	Description
informational	Usage of system physical capacity decreased to *current*%.

POOL_CREATE

Severity	Description

Severity	Description
informational	Storage Pool of size *pool.size*GB*sparse_type* was created with name '*pool.name*'.

POOL_CREATE_FAILED_TOO_MANY

Severity	Description	Troubleshooting
warning	Storage Pool with name '*name*' could not be created. You are attempting to add more Storage Pools than the system permits.	Delete Storage Pools to allow new ones to be created.

POOL_RENAME

Severity	Description
informational	Storage Pool with name '*old_name*' was renamed '*pool.name*'.

POOL_RESIZE

Severity	Description
informational	Storage Pool with name '*pool.name*' was resized from size *old_size*GB*old_sparse_type* to *pool.size*GB*sparse_type*.

POOL_RESIZE_SNAPSHOTS

Severity	Description
informational	Snapshot size of Storage Pool with name '*pool.name*' was resized from size *old_size*GB to *pool.snapshot_size*GB.

POOL_CHANGE_LOCK_BEHAVIOR

Severity	Description
informational	Lock Behavior of Storage Pool with name '*pool.name*' is now '*state*'.

POOL_CONVERTED_TO_SPARSE

Severity	Description
informational	Storage Pool with name '*pool.name*' was converted to sparse.

POOL_CONVERTED_TO_REGULAR

Severity	Description
informational	Storage Pool with name '*pool.name*' was converted to regular.

POOL_CHANGE_PERF_CLASS

Severity	Description
informational	Performance Class of Storage Pool with name '*pool.name*' is now '*pool.perf_class*'.

POOL_CONFIG_SNAPSHOTS

Severity	Description
informational	Management policy of Mirroring snapshots of Storage Pool with name '*pool.name*' has changed'.

POOL_DELETE

Severity	Description
informational	Storage Pool with name '*pool.name*' was deleted.

FLASH_VDISK_TOO_SMALL

Severity	Description
major	Flash vdisk *name* on *Flash Enclosure* is too small.

FLASH_VDISK_LARGER_THAN_EXPECTED

Severity	Description
warning	Flash vdisk *name* on *Flash Enclosure* is larger than expected.

DATA_REDUCTION_TIER_IS_OFFLINE

Severity	Description
critical	Data reduction tier moved to offline mode

DATA_REDUCTION_RECOVERY_STARTED

Severity	Description
Informational	Data reduction recovery process started

DATA_REDUCTION_RECOVERY_FINISHED

Severity	Description
Informational	Data reduction recovery process ended

DATA_REDUCTION_RECOVERY_FAILED

Severity	Description
Minor	Data reduction recovery process failed with reason: '*reason*'

DATA_REDUCTION_RECOVERY_ABORT_STARTED

Severity	Description
Informational	Data reduction recovery abort started

DATA_REDUCTION_RECOVERY_ABORT_FINISHED

Severity	Description
Informational	Data reduction recovery abort finished

DATA_REDUCTION_TIER_IS_ONLINE

Severity	Description
informational	Data reduction tier moved to online state

DATA_REDUCTION_TIER_BECOMING_OFFLINE

Severity	Description
major	Data reduction tier is transitioning to offline mode

DATA_REDUCTION_RESUME_ONLINE_START

Severity	Description
informational	Data reduction is resuming from offline mode

DATA_REDUCTION_RESUME_ONLINE_FAILED

Severity	Description
warning	Data reduction failed resuming from offline mode

DATA_REDUCTION_COMPRESSION_ADAPTER_FAILED

Severity	Description
Major	Data Reduction node *reporting_node* reported about compression adapter failure and will be killed.

FLASH_COMPONENT_INITIALIZING

Severity	Description	Troubleshooting
informational	*Flash Component ID* initializing.	Contact IBM Support

FLASH_COMPONENT_OK

Severity	Description	Troubleshooting
informational	*Flash Component ID* status ok.	Contact IBM Support

FLASH_COMPONENT_FAILED

Severity	Description	Troubleshooting
variable	*Flash Component ID* has failed.	Contact IBM Support

FLASH_COMPONENT_REPLACED

Severity	Description	Troubleshooting
informational	*Component ID* was replaced. New serial *New Serial* Old serial *Old Serial*	Contact IBM Support

FLASH_FW_HOT_UPGRADE_STARTED

Severity	Description
informational	Start upgrade to version *version*. Enclosure id *component id*

FLASH_FW_HOT_UPGRADE_FINISHED

Severity	Description
informational	Finished upgrade to version *version*. Enclosure id *component id*

FLASH_UPGRADE_RESUMED

Severity	Description
informational	Finished upgrade resumed. Enclosure id *component id*

FLASH_FW_HOT_UPGRADE_RESUMED

Severity	Description
informational	Finished upgrade resumed. Enclosure id *component id*

FLASH_UPGRADE_STOPPED

Severity	Description
major	Flash upgrade stopped after *percents*%, reason: *reason*. Enclosure id *component id*

FLASH_FW_HOT_UPGRADE_STOPPED

Severity	Description
major	Flash upgrade stopped after *percents*%, reason: *reason*. Enclosure id *component id*

FLASH_FW_HOT_UPGRADE_FAILED

Severity	Description
major	Flash firmware hot upgrade failed, reason: *reason*. error: *ccl_error*. Enclosure id *component id*

FLASH_FW_HOT_UPGRADE_TIMEOUT

Severity	Description
major	Timeout while upgrading *component id*, progress: *percents*

FLASH_RAID_STATUS_CHANGED

Severity	Description
variable	Flash array raid status changed to *status*. Enclosure id *component id*.

FLASH_ARRAY_STATUS_CHANGED

Severity	Description
variable	Flash array status changed to *status*. Enclosure id *component id*.

FLASH_ENCRYPTION_STATUS_CHANGED

Severity	Description
informational	Encryption enabled changed to *encrypted*. Enclosure id *component id*.

FLASH_ENCRYPTION_ENABLE_NOT_ALLOWED

Severity	Description
major	Encryption enable not allowed: *reason*. Enclosure id *component id*.

FLASH_ENCRYPTION_DISABLE_NOT_ALLOWED

Severity	Description
major	Encryption disable not allowed: *reason*. Enclosure id *component id*.

ENCRYPT_ENABLE_FLASH_ENCLOSURE_FAILED

Severity	Description
major	Encryption enable failed for *component id*: *reason*

ENCRYPT_DISABLE_FLASH_ENCLOSURE_FAILED

Severity	Description
major	Encryption disable failed for *component id*: *reason*

FLASH_ENCRYPTION_UNLOCK_FAILED

Severity	Description
major	Flash Encryption unlock failed. Enclosure: *component id*.

FLASH_ENCLOSURE_WIPEOUT_FAILED

Severity	Description
major	Wipeout failed for *component id*: *reason*

FLASH_BBU_CHARGING_STATUS_CHANGED

Severity	Description
informational	BBU charging status changed to *status*. BBU: *component id*.

FLASH_BBU_CALIBRATION_STARTED

Severity	Description
informational	BBU calibration started, BBU: *component id*.

FLASH_BBU_CALIBRATION_STOPPED

Severity	Description
informational	BBU calibration stopped, BBU: *component id*.

FLASH_BBU_CALIBRATION_FAILED

Severity	Description
minor	BBU calibration failed, BBU: *component id*.

FLASH_CANISTER_CONNECTED_VIA_SERIAL_CABLE

Severity	Description
informational	Established serial connection with *component id*.

FLASH_CANISTER_CONNECTION_VIA_SERIAL_OK

Severity	Description
informational	Established serial connection with *component id*.

NO_CONNECTION_TO_FLASH_CANISTER_VIA_SERIAL

Severity	Description	Troubleshooting
warning	Failed to connect to *component id* via serial.	Contact IBM Support

FLASH_CANISTER_NO_CONNECTION_VIA_SERIAL

Severity	Description	Troubleshooting
warning	Failed to connect to *component id* via serial.	Contact IBM Support

NO_CONNECTION_TO_FLASH_CANISTER_VIA_ETH

Severity	Description	Troubleshooting
warning	Failed to connect to *component id* via ethernet.	Contact IBM Support

FLASH_CANISTER_NO_CONNECTION_VIA_ETH

Severity	Description	Troubleshooting
warning	Failed to connect to *component id* via ethernet.	Contact IBM Support

FLASH_CANISTER_CONNECTION_VIA_ETH_OK

Severity	Description	Troubleshooting
informational	Established connection to *component id* via ethernet.	Contact IBM Support

FLASH_CANISTER_IP_CHANGED

Severity	Description	Troubleshooting
informational	Changed connected ip of *component id* to *IP*.	Contact IBM Support

NO_CONNECTION_TO_FLASH_CANISTER_VIA_GW

Severity	Description	Troubleshooting
warning	Failed to connect to *component id* via *module id*.	Contact IBM Support

FLASH_CANISTER_NO_CONNECTION_VIA_GW

Severity	Description	Troubleshooting
warning	Failed to connect to *component id* via *module id*.	Contact IBM Support

FLASH_CANISTER_CONNECTED_VIA_GATEWAY

Severity	Description	Troubleshooting
informational	Established connection to *component id* via *module id*.	Contact IBM Support

FLASH_CANISTER_CONNECTION_VIA_GW_OK

Severity	Description	Troubleshooting
informational	Established connection to *component id* via *module id*.	Contact IBM Support

FLASH_CANISTER_ETH_LINK_MISWIRE

Severity	Description	Troubleshooting
minor	Flash canister *component id* is miswired.	Contact IBM Support

FLASH_CANISTER_ETH_LINK_WIRING_OK

Severity	Description	Troubleshooting
informational	Flash canister *component id* is no longer miswired.	Contact IBM Support

FLASH_CANISTER_GET_CONF_FAILED

Severity	Description	Troubleshooting
minor	Failed to get the configuration from *component id*.	Contact IBM Support

FLASH_CANISTER_GET_CONF_OK

Severity	Description	Troubleshooting
informational	Succeeded to get the configuration from *component id*.	Contact IBM Support

FLASH_CANISTER_IN_SERVICE_MODE

Severity	Description	Troubleshooting
major	Flash canister is in service mode *service*, canister: *component id*.	Contact IBM Support

FLASH_CANISTER_NO_LONGER_IN_SERVICE_MODE

Severity	Description
informational	Flash canister is no longer in service mode: *component id*.

FLASH_COMPONENT_TEMPERATURE_OK

Severity	Description
informational	Flash component *component id* temperature is within allowed limits.

FLASH_COMPONENT_TEMPERATURE_ABOVE_NORMAL

Severity	Description
warning	Flash component *component id* temperature is above normal.

FLASH_COMPONENT_TEMPERATURE_HIGH

Severity	Description
major	Flash component *component id* temperature is high, it exceeds operational level.

FLASH_COMPONENT_TEMPERATURE_CRITICALLY_HIGH

Severity	Description
critical	Flash component *component id* temperature is critical.

FLASH_ENCLOSURE_STATUS_CHANGED

Severity	Description	Troubleshooting
informational	*component id* status changed to *New Status*.	Contact IBM Support

FLASH_PSU_FAN_FAILED

Severity	Description
minor	*component id* fan has failed.

FLASH_PSU_FAN_OK

Severity	Description
informational	*component id* is ok.

FLASH_PSU_HAS_NO_INPUT_POWER

Severity	Description
major	PSU power supply has no input (AC) power. PSU id *component id*.

FLASH_PSU_HAS_INPUT_POWER

Severity	Description
informational	PSU power supply has input (AC) power. PSU id *component id*.

FLASH_ENCLOSURE_VERSION_IS_UNEXPECTED

Severity	Description
major	*component id* version is *version*, expected version is *expected*.

FLASH_ENCLOSURE_NEWER_VERSION_EXISTS

Severity	Description
minor	*component id* version *version* is supported. newer version *expected* exists.

FLASH_PSU_COMMUNICATION_ERROR

Severity	Description
minor	PSU communication error. PSU id *component id*.

FLASH_BBU_VPD_IS_NOT_VALID

Severity	Description
major	BBU VPD is not valid. BBU id *component id*.

FLASH_CARD_UNSUPPORTED

Severity	Description
warning	Flash card is unsupported. Flash card id *component id*.

FLASH_BBU_END_OF_LIFE

Severity	Description
warning	Battery is at end of life. BBU id *component id*.

FLASH_BBU_NEARING_END_OF_LIFE

Severity	Description
warning	Battery is nearing end of life. BBU id *component id*.

FLASH_CARD_COMMUNICATION_ERROR

Severity	Description
minor	Flash card communication error. Flash card id *component id*.

FLASH_FAN_COMMUNICATION_ERROR

Severity	Description
minor	Fan communication error. Fan: *component id*.

FLASH_ENCLOSURE_THERMAL_THREASHOLD_EXCEEDED

Severity	Description
critical	Flash enclosure *component id* has shutdown after exceeding the thermal threshold.

FLASH_ENCLOSURE_ARRAY_OFFLINE

Severity	Description
critical	Flash enclosure *component id* array is offline.

FLASH_ENCLOSURE_STARTED_PHASEOUT

Severity	Description
informational	System started phasing out *Component ID*.

FLASH_ENCLOSURE_FINISHED_PHASEOUT

Severity	Description
informational	System finished phasing out *Component ID*.

FLASH_ENCLOSURE_STARTED_PHASEIN

Severity	Description
informational	System started phasing in *Component ID*.

FLASH_ENCLOSURE_FINISHED_PHASEIN

Severity	Description
informational	System finished phasing in *Component ID*.

FLASH_CR_KEY_SETUP_FAILED

Severity	Description
major	Failed to set challenge-response key on '*Component ID*'.

FLASH_CR_KEY_SETUP_OK

Severity	Description
informational	Challenge-response key was successfully set on '*Component ID*'.

FLASH_CR_KEY_SETUP_STARTED

Severity	Description
informational	Challenge-response key set started on '*Component ID*'.

FC_PORT_HAS_FAILED

Severity	Description	Troubleshooting
major	FC port service *port* has failed due to *code* (attempt number *Number of retries*)	Contact IBM Support

NTP_SERVER_TIME_DIFFERENCE_TOO_BIG

Severity	Description	Troubleshooting
warning	NTP server *NTP Server* sent a transaction with time difference of *Delta* seconds which exceeds the maximal difference of *Max Allowed* seconds. Transaction will be ignored, please check NTP server's and system's times.	Please contact your Administrator.

IPSEC_TUNNEL_OPENED

Severity	Description
informational	The IPSec tunnel named '*name*' between module *Module* and *Right IP* was opened

IPSEC_TUNNEL_CLOSED

Severity	Description
informational	The IPSec tunnel named '*name*' between module *Module* and *Right IP* was closed

IP_ACCESS_CANNOT_RESOLVE_ADDRESS

Severity	Description
informational	Cannot resolve address '*address*' added to the IP access group *IP access group name*.

IP_ACCESS_FAILED_SETTING_RULES

Severity	Description
informational	Failed setting IP access rules.

USB_ETHERNET_INTERFACE_OK

Severity	Description
informational	USB ethernet interface on module *Module* was reset successfully and is now OK.

USB_ETHERNET_INTERFACE_IS_STILL_DOWN

Severity	Description	Troubleshooting
minor	USB ethernet interface on module *Module* is down and failed to reset.	Contact IBM Support

USB_ETHERNET_INTERFACE_FAILED

Severity	Description	Troubleshooting
warning	USB ethernet interface on module *Module* failed.	Contact IBM Support

MIRROR_CREATE

Severity	Description
informational	A remote mirror was defined for Volume '*local volume name*'on Target '*target name*'. Remote Volume is '*remote volume name*'.

CG_MIRROR_CREATE

Severity	Description
informational	A remote mirror was defined for Consistency Group '*local CG name*' on Target '*target name*'. Remote Consistency Group is '*remote CG name*'.

MIRROR_CREATE_SLAVE

Severity	Description
informational	A remote mirror was defined by Target '*target name*' for Volume '*local volume name*'. Remote Volume is '*remote volume name*'.

CG_MIRROR_CREATE_SLAVE

Severity	Description
informational	A remote mirror was defined by Target '*target name*' for CG '*local CG name*'. Remote CG is '*remote CG name*'.

MIRROR_SCHEDULE_CHANGE

Severity	Description
informational	Schedule of remote mirror of '*local peer name*' is now '*schedule name*'.

MIRROR_CREATE_FAILED_TARGET_NOT_CONNECTED

Severity	Description	Troubleshooting
warning	Target could not be reached. Target with name '*target.name*' is currently not connected.	Connect the target system to this system.

REMOTE_OPERATION_FAILED_TIMED_OUT

Severity	Description	Troubleshooting
warning	Operation on remote machine timed out. Invoking '*Function Name*' on target '*Target Name*' timed out.	Retry operation. If problem persists contact IBM support.

MIRROR_RESYNC_FAILED

Severity	Description	Troubleshooting
major	Synchronization of meta data with mirror failed. Configuration of remote mirror of volume '*local volume name*' on target '*target name*' does not match local configuration.	Make sure configuration on both machines is compatible and activate the mirror. If problem persists contact IBM support.

MIRROR_RESYNC_FAILED_DUE_TO_THIN_PROVISIONING

Severity	Description	Troubleshooting
major	Synchronization of bitmaps with mirror failed. Not enough hard capacity left in Pool of volume '*mirror.local_volume_name*'.	Delete unnecessary volumes in pool or enlarge the pool's hard size.

MIRROR_SYNC_STARTED

Severity	Description
informational	Synchronization of remote mirror of volume '*local volume name*' on Target '*target name*' has started.

MIRROR_SYNC_ENDED

Severity	Description
informational	Synchronization of remote mirror of peer '*local peer name*' on target '*target name*' has ended.

MIRROR_CANNOT_CREATE_SYNC_JOB_TOO_MANY_VOLUMES

Severity	Description
major	Synchronization of remote mirror of peer '*local peer name*' on target '*target name*' can not be synced , insufficent volume available for this operation.

MIRROR_CANNOT_CREATE_LRS_TOO_MANY_VOLUMES

Severity	Description
major	Synchronization of remote mirror of peer '*local peer name*' on target '*target name*' can not be synced , insufficent volume available for this operation.

MIRROR_REESTABLISH_FAILED_CONFIGURATION_ERROR

Severity	Description	Troubleshooting
major	Mirror reestablish failed. Either configuration of remote mirror of peer '*local peer name*' on target '*target name*' does not match local configuration.	Make sure configuration on both machines is compatible and activate the mirror. If problem persists contact IBM support.

MIRROR_ACTIVATE

Severity	Description
informational	The Remote Mirror of peer '*local peer name*' on Target '*target name*' was activated.

MIRROR_DEACTIVATE

Severity	Description
informational	The Remote Mirror of peer '*local peer name*' on Target '*target name*' was deactivated.

MIRROR_SLAVE_ACTIVATE

Severity	Description
informational	The mirror of peer '*local peer name*' on target '*target name*' was activated.

MIRROR_SLAVE_DEACTIVATE

Severity	Description
informational	The mirror of peer '*local peer name*' on target '*target name*' was deactivated.

MIRROR_DEACTIVATE_SECONDARY_LOCKED

Severity	Description
minor	The Remote Mirror of peer '*local peer name*' on Target '*target name*' was deactivated since the Pool on the secondary machine was locked.

MIRROR_DEACTIVATE_CONFIGURATION_ERROR

Severity	Description
minor	The Remote Mirror of peer '*local peer name*' on Target '*target name*' was deactivated since the Mirror configuration on the slave machine has changed.

MIRROR_DELETE

Severity	Description
informational	The Remote Mirror relation of peer '*local peer name*' to a peer on Target '*target name*' was deleted.

MIRROR_REVERSE_ROLE_TO_SLAVE

Severity	Description
informational	Local peer '*local peer name*' is now Slave of a peer on Target '*target name*'.

MIRROR_REVERSE_ROLE_TO_MASTER

Severity	Description
informational	Local peer '*local peer name*' is now Master of a peer on Target '*target name*'.

MIRROR_REVERSE_ROLE_OF_PEER_WITH_LCS_TO_MASTER

Severity	Description
informational	Local peer '*local peer name*' is now Master of a peer on Target '*target name*' external last consistent snapshot should be deleted manually .

MIRROR_SWITCH_ROLES_TO_SLAVE

Severity	Description

Severity	Description
informational	Local peer '*local peer name*' switched roles with peer on Target '*target name*'. It is now Slave.

MIRROR_SWITCH_ROLES_TO_MASTER

Severity	Description
informational	Local peer '*local peer name*' switched roles with peer on Target '*target name*'. It is now Master.

MIRROR_REESTABLISH_FAILED_TOO_MANY_VOLUMES

Severity	Description	Troubleshooting
major	Last Consistent Snapshot of Slave peer '*local peer name*' could not be created. Maximal number of Volumes are already defined.	Delete Volumes to allow new ones to be created. Activate Mirror on the Master Machine.

MIRROR_END_SYNC_FAILED_CONFIGURATION_ERROR

Severity	Description	Troubleshooting
major	Configuration of remote mirror of peer '*local peer name*' on target '*target name*' does not match local configuration.	Make sure configuration on both machines is compatible and activate the mirror. If problem persists contact IBM support.

MIRROR_CHANGE_DESIGNATION

Severity	Description
informational	Local peer '*local peer name*' switched its designated role with peer on Target '*target name*'. It is now *designation*.

MIRROR_CANCEL_SNAPSHOT

Severity	Description
informational	All mirrored snapshots which were created for Mirror of peer '*local peer name*' and were not yet synchronized will not be mirrored in the remote machine.

DM_DEFINE

Severity	Description
informational	Data Migration was defined to Volume '*local volume name*' from Target '*target name*'.

DM_SYNC_STARTED

Severity	Description
informational	Migration to volume '*local volume name*' from Target '*target name*' has started.

DM_SYNC_ENDED

Severity	Description
informational	Migration to volume '*local volume name*' from target '*target name*' is complete.

DM_SYNC_ENDED_WITH_ERRORS

Severity	Description
Critical	Migration to volume '*local volume name*' from target '*target name*' has completed with *medium_errors_in_data_migration* error(s). Check previous events related to this volume for the list of affected LBAs.'.

DM_ACTIVATE

Severity	Description
informational	Migration to Volume '*local volume name*' from Target '*target name*' was activated.

DM_DEACTIVATE

Severity	Description
informational	Migration to Volume '*local volume name*' from Target '*target name*' was deactivated.

DM_DEACTIVATE_LUN_UNAVAILABLE

Severity	Description
minor	Migration to Volume '*local volume name*' from Target '*target name*' was deactivated since LUN is not available on one of the active paths to the target.

DM_DELETE

Severity	Description
informational	Definition of Data Migration to Volume '*local volume name*' from Target '*target name*' was deleted.

SCHEDULE_CREATE

Severity	Description
informational	Schedule was created with name '*schedule name*'.

SCHEDULE_UPDATE

Severity	Description
informational	Schedule with name '*schedule name*' was updated.

SCHEDULE_RENAME

Severity	Description
informational	Schedule with name '*old_name*' was renamed '*schedule name*'.

SCHEDULE_DELETE

Severity	Description
informational	Schedule with name '*schedule name*' was deleted.

MIRROR_RPO_OK

Severity	Description
informational	Mirror of local peer '*local peer name*' is now ahead of its specified RPO.

MIRROR_RPO_LAGGING

Severity	Description
informational	Mirror of local peer '*local peer name*' is now behind its specified RPO.

MIRROR_CHANGE_RPO

Severity	Description
informational	RPO or Mirror of local peer '*local peer name*' is now *RPO*.

MIRROR_IS_LAGGING_BEYOND_PERCENT_THRESHOLD

Severity	Description
warning	Last Replication Time of Mirror of local peer '*local peer name*' is *Last Replication Time*.

MIRROR_AUTO_FIX_REACHED_LIMIT

Severity	Description
warning	A remote checksum diff for mirror '*local peer name*' cannot be fixed automatically because we reached the auto fix limit.

MIRROR_IS_LAGGING_BEYOND_ABSOLUTE_THRESHOLD

Severity	Description
warning	Last Replication Time of Mirror of local peer '*local peer name*' is *Last Replication Time*.

MIRROR_INCOMPATIBLE_VERSION_FOR_UNMAP_SUPPORT

Severity	Description
warning	Mirror of peer '*local peer name*' on target '*target name*' cannot support unmap, remote machine has incompatible version.

MULTISITE_DEFINE

Severity	Description
informational	A Multi-site Master was defined for volume '*local volume name*' on SMaster target and vol '*SMaster target name*' '*SMaster volume name*', and Slave target and vol '*slave target name*' '*slave volume name*'.

MULTISITE_DEFINE_SLAVE

Severity	Description
informational	A Multi-site Slave was defined for volume '*local volume name*' on Master target and vol '*master target name*' '*master volume name*', and SMaster target and vol '*SMaster target name*' '*SMaster volume name*'.

MULTISITE_DEFINE_SMASTER

Severity	Description
informational	A Multi-site SMaster was defined for volume '*local volume name*' on Master target and vol '*master target name*' '*master volume name*', and Slave target and vol '*slave target name*' '*slave volume name*'.

CG_MULTISITE_DEFINE

Severity	Description
informational	A Multi-site Master was defined for consistency group '*local CG name*' on SMaster target and CG '*SMaster target name* '*SMaster CG name*', and Slave target and vol '*slave target name*' '*slave CG name*.

CG_MULTISITE_DEFINE_SLAVE

Severity	Description
informational	A Multi-site Slave was defined for consistency group '*local CG name*' on Master target and CG '*master target name*' '*master CG name*', and SMaster target and CG '*SMaster target name*' '*SMaster CG name*'.

CG_MULTISITE_DEFINE_SMASTER

Severity	Description
informational	A Multi-site SMaster was defined for consistency group '*local CG name*' on Master target and vol '*master target name*' '*master CG name*', and Slave target and CG '*slave target name*' '*slave CG name*'.

MULTISITE_DELETE

Severity	Description
informational	The Multi-site relation of peer '*local peer name*' was deleted.

MULTISITE_CHANGE_SLAVE_ROLE_TO_MASTER

Severity	Description
informational	The Multi-site relation of peer '*local peer name*' was changed to standalone.

MULTISITE_CHANGE_MASTER_ROLE_TO_SLAVE

Severity	Description
informational	The Multi-site relation of peer '*local peer name*' was changed to Slave.

MULTISITE_CHANGE_SMASTER_ROLE_TO_MASTER

Severity	Description
informational	The Multi-site relation of peer '*local peer name*' was changed to Master.

MULTISITE_CHANGE_MASTER_ROLE_TO_SMASTER

Severity	Description
informational	The Multi-site relation of peer '*local peer name*' was changed to SMaster.

MULTISITE_STANDBY_MIRROR_REGISTERED

Severity	Description
informational	The Multi-site relation of peer '*local peer name*' registered a standby mirror on SMaster system

MULTISITE_COMPROMISED

Severity	Description
warning	The state of the Multi-site relation on peer '*local peer name*' changed to compromised, reason: *Compromise Reason*

MULTISITE_RESTORED

Severity	Description
informational	The state of the Multi-site relation on peer '*local peer name*' changed from compromised to operational or degraded

MULTISITE_SWITCH_ROLES_TO_MASTER

Severity	Description
informational	Local peer '*local peer name*' is now Master of Multi-site with SMaster target '*SMaster target name*', Slave target '*slave target name*'.

MULTISITE_SWITCH_ROLES_TO_SMASTER

Severity	Description
informational	Local peer '*local peer name*' is now SMaster of Multi-site with Master target '*master target name*', Slave target '*slave target name*'.

MULTISITE_AUTOMATIC_ACTIVATION_OF_STANDBY_ASYNC_FAILED

Severity	Description
major	Automatic activation of the standby Async mirror relation '*slave volume name*' had failed .

MAP_VOLUME

Severity	Description
informational	Volume with name '*volume.name*' was mapped to LUN '*LUN*' for *host_or_cluster* with name '*host*'.

MAP_PROXY_VOLUME

Severity	Description
informational	IBM Hyper-Scale Mobility Volume with name '*name*' was mapped to LUN '*LUN*' for *host_or_cluster* with name '*host*'.

UNMAP_VOLUME

Severity	Description
informational	Volume with name '*volume.name*' was unmapped from *host_or_cluster* with name '*host*'.

UNMAP_PROXY_VOLUME

Severity	Description
informational	IBM Hyper-Scale Mobility Volume with name '*name*' was unmapped from *host_or_cluster* with name '*host*'.

SPECIAL_TYPE_SET

Severity	Description
informational	Type of *host_or_cluster* with name '*host*' was set to '*type*'.

SERVICE_HAS_FAILED

Severity	Description	Troubleshooting
variable	*Component ID* has failed.	Contact IBM Support

SERVICE_FAILED_TO_PHASEIN

Severity	Description	Troubleshooting
major	*Component ID* failed to phase-in.	Contact IBM Support

SERVICE_FAILED_TO_RESTART

Severity	Description	Troubleshooting
major	*Component ID* failed to restart.	Contact IBM Support

MODULE_FAILED

Severity	Description	Troubleshooting
critical	*Component ID* failed.	Contact IBM Support

NODE_OK

Severity	Description
informational	Node #*Node ID* of type *Node Type* on *Module Component ID* has been revived and OK now. The node pid: *Process ID*

NODE_AUTO_RESET_FAILED

Severity	Description	Troubleshooting
minor	Node #*Node ID* of type *Node Type* on *Module Component ID* has failed and failed to auto-revive.	Contact IBM Support

DISK_HAS_FAILED

Severity	Description	Troubleshooting
variable	Disk *Component ID* Failed.	Please contact your Administrator.

SSD_HAS_FAILED

Severity	Description	Troubleshooting
major	SSD *Component ID* Failed.	Please contact your Administrator.

VAULT_DEVICE_HAS_FAILED

Severity	Description	Troubleshooting
minor	Vault device *Component ID* Failed.	Please contact your Administrator.

COMPONENT_TEST_OF_DISK_HAS_FAILED

Severity	Description	Troubleshooting
major	Test of *Component ID* has failed with error *Error*.	Please contact your Administrator.

COMPONENT_TEST_OF_SSD_HAS_FAILED

Severity	Description	Troubleshooting
major	Test of *Component ID* has failed with error *Error*.	Please contact your Administrator.

COMPONENT_TEST_OF_BOOT_MEDIA_HAS_FAILED

Severity	Description	Troubleshooting
major	Test of *Component ID* has failed with error *Error*.	Please contact your Administrator.

COMPONENT_TEST_OF_VAULT_DEVICE_HAS_FAILED

Severity	Description	Troubleshooting
major	Test of *Component ID* has failed with error *Error*.	Please contact your Administrator.

BOOT_MEDIA_COMPONENT_TEST_FAILED

Severity	Description	Troubleshooting
major	Test of *Component ID* has failed with error *Error*.	Please contact your Administrator.

DISK_STARTED_PHASEOUT

Severity	Description
informational	System started phasing out *Component ID*.

DISK_STARTED_PHASEIN

Severity	Description
informational	System started phasing in *Component ID*.

DISK_FINISHED_PHASEIN

Severity	Description
informational	System finished phasing in *Component ID*.

DISK_FINISHED_PHASEOUT

Severity	Description
informational	System finished phasing out *Component ID*.

DISK_RECOVERED

Severity	Description
critical	Disk *Component ID* is functioning again.

MODULE_STARTED_PHASEOUT

Severity	Description
informational	System started phasing out *Component ID*.

MODULE_FINISHED_PHASEOUT

Severity	Description
informational	System finished phasing out *Component ID*.

MODULE_STOPPED_PHASEOUT_DUE_TO_MANAGEMENT_REQUIREMENT

Severity	Description
major	System stopped phasing out *Component ID* due to management requirement.

START_WORK

Severity	Description
informational	System has entered ON state.

SYSTEM_HAS_ENTERED_MAINTENANCE_MODE

Severity	Description
warning	System has entered MAINTENANCE state [*Reason*]

SYSTEM_LEFT_CHARGING_STATE

Severity	Description
informational	System is sufficiently charged.

USER_SHUTDOWN

Severity	Description
major	System is shutting down due to a user request.

EMERGENCY_SHUTDOWN_NOW

Severity	Description	Troubleshooting
critical	System is shutting down in emergency shutdown mode due to: *Emergency Shutdown Reason*.	Please contact your Administrator.

SHOULD_HAVE_BEEN_EMERGENCY_SHUTDOWN

Severity	Description
critical	An emergency shutdown has been detected, but emergency shutdown is disabled for the detected reason. Shutdown reason: *Shutdown Reason*.

DATA_SERVICE_STARTED_PHASEOUT

Severity	Description

Severity	Description
informational	System started phasing out *Component ID*.

DATA_SERVICE_STARTED_PHASEIN

Severity	Description
informational	System started phasing in *Component ID*.

DATA_SERVICE_FINISHED_PHASEIN

Severity	Description
informational	System finished phasing in *Component ID*.

DATA_SERVICE_FINISHED_PHASEOUT

Severity	Description
informational	System finished phasing out *Component ID*.

TXN_REBUILD_STARTED

Severity	Description
informational	Start rebuild process of txns.

TXN_REBUILD_ENDED

Severity	Description
informational	End rebuild process for txns.

TXN_REDIST_STARTED

Severity	Description
informational	Start redist process for txns.

TXN_REDIST_ENDED

Severity	Description
informational	End redist process for txns.

DISK_MARKED_TO_PHASEOUT

Severity	Description
informational	System started phasing out *Component ID*.

DISK_MARKED_TO_PHASEIN

Severity	Description
informational	System started phasing out *Component ID*.

CANNOT_CREATE_NEW_DATA_DISTRIBUTION

Severity	Description
informational	System cannot phaseout disks for storage medium *Storage Medium*.

CANNOT_RESIZE_FLASH_MEDIUM_POOLS

Severity	Description
informational	System failed to resize flash medium pools for TMS phaseout.

DATA_PROTECTION_STATUS_CHANGED

Severity	Description
variable	Data protection status has changed from '*old_status*' to '*new_status*'

VAULT_DEVICE_SECURE_ERASE_PROCESS_SUCCESSFUL

Severity	Description
informational	Vault device secure erase process successful.

VAULT_DEVICE_SECURE_ERASE_PROCESS_FAILED

Severity	Description
major	Vault device secure erase process failed [*Reason*]

VAULT_DEVICE_FAILED_SECURE_ERASE

Severity	Description
major	Secure erase for *Component ID* failed. [*Reason*].

SYSTEM_PHYSICAL_CAPACITY_CHANGED

Severity	Description
informational	System physical capacity is now *Capacity*GB.

SYSTEM_EFFECTIVE_CAPACITY_CHANGED

Severity	Description
informational	System effective capacity is now *Capacity*GB.

SYSTEM_OUT_OF_PHYSICAL_SPACE

Severity	Description
critical	System has run out of physical capacity. All volumes are now write-protected.

SYSTEM_NORMAL_OPERATION_RESUMED

Severity	Description
informational	Normal operation is resumed. Volumes have been restored to their original write-protection state.

SYSTEM_FREE_PHYSICAL_CAPACITY_REACHED_RESERVED_LIMIT

Severity	Description
critical	System physical capacity usage reached reserved limit. All volumes are now write-protected.

ENCRYPT_ENABLE_STARTED

Severity	Description
informational	Starting encryption activation. This process can take several minutes to complete.

ENCRYPT_ENABLE_COMPLETED

Severity	Description
informational	Encryption is in effect.

ENCRYPT_ENABLE_NOT_COMPLETED

Severity	Description	Troubleshooting
major	Cannot complete encryption activation because *reason*. *Count* vault device(s) and *Count* flash enclosure(s) could not be enrolled.	Please contact technical support

ENCRYPT_DISABLE_STARTED

Severity	Description
informational	Starting encryption deactivation. This process can take several minutes to complete.

ENCRYPT_DISABLE_COMPLETED

Severity	Description

Severity	Description
informational	Encryption is no longer in effect.

ENCRYPT_DISABLE_NOT_COMPLETED

Severity	Description	Troubleshooting
major	Cannot complete encryption deactivation because *reason*. *Count* vault device(s) and *Count* flash enclosure(s) could not be crypto erased.	Please contact technical support

ENCRYPT_KEYSERVER_ADDED

Severity	Description
informational	A key server named '*Key Server Name*' was added.

ENCRYPT_KEYSERVER_DELETED

Severity	Description
informational	Key server '*Key Server Name*' was deleted.

ENCRYPT_KEYSERVER_EDITED

Severity	Description
informational	Details of key server '*Key Server Name*' were modified.

ENCRYPT_KEYSERVER_RENAMED

Severity	Description
informational	Key server '*Old Name*' was renamed to '*New Name*'.

ENCRYPT_KEYSERVER_CHECK_STATUS_STARTED

Severity	Description
informational	Start checking connectivity status of the keyservers currently defined in the system. This process can take several minutes to complete.

ENCRYPT_KEYSERVER_CHECK_STATUS_COMPLETED

Severity	Description
informational	Completed checking connectivity status of the keyservers currently defined in the system.

ENCRYPT_KEYSERVER_REKEY_COMPLETED

Severity	Description
informational	Key server '*Key Server Name*' rekey completed.

ENCRYPT_LOCAL_REKEY_COMPLETED

Severity	Description
informational	Local key rekey completed.

ENCRYPT_CHANGE_KEY_SCHEME_COMPLETED

Severity	Description
informational	Change key scheme from external to local key completed.

ENCRYPT_CHANGE_KEY_SCHEME_FAILED

Severity	Description
major	Change key scheme from external to local key failed because *failure reason*.

ENCRYPT_CHANGE_KEY_SCHEME_ROLLBACK_FAILED

Severity	Description
major	Cannot rollback change key scheme.

ENCRYPT_UNABLE_TO_UPDATE_KEY_DURING_DEACTIVATE_ON_KEYSERVER

Severity	Description
major	Could not update key server '*Keyserver Name*' regarding encryption deactivation. Please check key server status.

ENCRYPT_KEYSERVER_REKEY_FAILED

Severity	Description
major	Cannot complete rekey with key server '*Key Server Name*'.

ENCRYPT_LOCAL_REKEY_FAILED

Severity	Description
major	Local key rekey failed because *failure reason*.

ENCRYPT_KEYSERVER_REKEY_ROLLBACK_FAILED

Severity	Description
major	Cannot rollback failed rekey with key server '*Key Server Name*'.

ENCRYPT_LOCAL_REKEY_ROLLBACK_FAILED

Severity	Description
major	Cannot rollback failed local key rekey.

ENCRYPT_RECOVERY_KEY_ENTERED

Severity	Description
informational	Valid recovery key share was entered by user '*User Name*'.

ENCRYPT_INVALID_RECOVERY_KEY_ENTERED

Severity	Description
major	Invalid recovery key share was entered by user '*User Name*'.

ENCRYPT_RECOVERY_KEYS_GENERATED

Severity	Description
informational	Recovery keys created.

ENCRYPT_RECOVERY_KEY_REKEY_SUCCESS

Severity	Description
informational	Recovery key rekey was successful.

ENCRYPT_RECOVERY_KEY_REKEY_FAIL

Severity	Description
major	Recovery key rekey failed.

ENCRYPT_RECOVERY_KEY_VERIFIED

Severity	Description
informational	Recovery key verified successfully for user '*User Name*'.

ENCRYPT_RECOVERY_KEY_VERIFY_FAILED

Severity	Description
major	Recovery key verification failed for user '*User Name*'.

ENCRYPT_RECOVERY_KEY_ALL_SHARES_VERIFIED

Severity	Description
informational	All recovery key shares have been verified.

ENCRYPT_KR_WRITE_FAILED

Severity	Description	Troubleshooting
critical	Key repository write failed with error code: *rc*.	Please contact technical support

ENCRYPT_KR_READ_FAILED

Severity	Description	Troubleshooting
major	Key repository read failed with error code: *rc*.	Please contact technical support

ENCRYPT_UNABLE_TO_RETRIEVE_KEY_FROM_KEYSERVER

Severity	Description
major	Failed to retrieve key from key server '*Keyserver Name*' via *TEXT* on module *node id*. Please verify that the key server type and version are supported. If so, please check its status.

ENCRYPT_RECOVERY_KEY_RECOVER_SUCCESSFUL

Severity	Description
informational	Key recovery was successful, unlocking system.

ENCRYPTION_CERTIFICATE_FOR_XIV_IS_NOT_INSTALLED

Severity	Description	Troubleshooting
critical	XIV certificate is not installed.	Check output of pki_list for a certificate named XIV and contact technical support

ENCRYPT_UNABLE_TO_DELETE_MASTER_KEYSERVER

Severity	Description
informational	Deletion of master key server '*Keyserver Name*' is not allowed. Please define another key server as master first'.

ENCRYPTION_SKMIP_ERROR

Severity	Description	Troubleshooting
major	Module *Module* reported *Keyserver Name* returned error:*error code - TEXT*	Please contact the next level of support.

WIPEOUT_STARTED

Severity	Description
informational	Starting the wipeout process. This process may take several minutes to complete.

WIPEOUT_COMPLETED

Severity	Description
informational	The wipeout process finished successfully.

WIPEOUT_NOT_COMPLETED

Severity	Description	Troubleshooting
major	Cannot complete the wipeout process because *reason*. *Count* vault device(s) and *Count* flash enclosure(s) could not be crypto erased.	Contact IBM Support

DIMM_FAILED

Severity	Description	Troubleshooting
major	*Component ID* has failed. Hardware status: *Status*.	Contact IBM Support

CPU_FAILED

Severity	Description	Troubleshooting
major	*Component ID* has failed. Hardware status: *Status*.	Contact IBM Support

NIC_FAILED

Severity	Description	Troubleshooting
major	*Component ID* has failed. Hardware status: *Status*.	Contact IBM Support

MODULE_BBU_FAILED

Severity	Description	Troubleshooting
major	*BBU id* has failed. Hardware status: '*Status*'. BBU state: '*State*'.	Contact IBM Support

MODULE_BBU_OK

Severity	Description
informational	*BBU id* is now OK.

DIMM_WAS_REMOVED

Severity	Description	Troubleshooting
minor	The DIMM with serial number '*Serial*' was removed from *ModuleId*.	Was this DIMM actually removed?

PSU_CHANGE_DETECTED

Severity	Description	Troubleshooting
informational	*Component ID* has been changed from a serial number '*old_serial*', part number '*old_part_number*', to serial number '*new_serial*' and part number '*new_part_number*'.	Was this PSU actually replaced?

PSU_WAS_REMOVED

Severity	Description	Troubleshooting
warning	*Component ID* with a serial number '*Serial*' and part number '*Part Number*' was removed from the system.	Was this PSU actually removed?

PSU_MISSING_FROM_INSTALL

Severity	Description	Troubleshooting
major	System was installed without *Component ID*.	This PSU was not in the system at install time.

PSU_WAS_INSTALLED

Severity	Description	Troubleshooting
informational	*Component ID* with a serial number '*Serial*' and part number '*Part Number*' was installed in the system.	Was this PSU actually installed?

NIC_CHANGE_DETECTED

Severity	Description	Troubleshooting
major	*Component ID* has been changed from a serial of *old_serial* to *new_serial*.	Was this NIC actually replaced?

MODULE_BBU_IS_TOO_OLD

Severity	Description
major	*BBU id* installed more than *max use time* months after manufacturing date *manufacturing date*.

MODULE_BBU_SHELF_TIME_EXCEEDED

Severity	Description
major	*BBU id* installed more than *max shelf time* months after manufacturing date *manufacturing date*.

TECHNICIAN_WORK_STARTED

Severity	Description
informational	Technician work has started, expected to end at *End Time*. Comment: *Comment*.

TECHNICIAN_WORK_ENDED

Severity	Description
informational	Technician work has ended after *Elapsed Time* minutes. Comment: *Comment*.

TECHNICIAN_WORK_TIMED_OUT

Severity	Description
warning	Technician work has timed out after *Elapsed Time* minutes. Comment: *Comment*.

XIV_SUPPORT_ENABLED

Severity	Description
informational	XIV support access from *From* is enabled from *Start Time* until *Finish Time*. Comment: *Comment*.

XIV_SUPPORT_ENABLED_NO_TIME_LIMIT

Severity	Description
informational	XIV support access from *From* is enabled from *Start Time* until explicitly disabled. Comment: *Comment*.

XIV_SUPPORT_DISABLED

Severity	Description
informational	XIV support access is disabled.

XIV_SUPPORT_WINDOW_TIMEOUT

Severity	Description
informational	XIV support work window timeout is expired.

HOST_DEFINE

Severity	Description
informational	Host of type *host.type* was defined with name '*host.name*'.

HOST_UPDATE

Severity	Description
informational	Host named '*host.name*' was updated.

CLUSTER_CREATE

Severity	Description
informational	Cluster was defined with name '*cluster.name*'.

HOST_DEFINE_FAILED_TOO_MANY

Severity	Description	Troubleshooting
warning	Host with name '*name*' could not be defined. You are attempting to define more hosts than the system permits.	Delete Hosts to allow new ones to be defined.

CLUSTER_CREATE_FAILED_TOO_MANY

Severity	Description	Troubleshooting
warning	Cluster with name '*name*' could not be defined. You are attempting to define more Clusters than the system permits.	Delete Clusters to allow new ones to be defined.

HOST_RENAME

Severity	Description
informational	Host with name '*old_name*' was renamed '*host.name*'.

CLUSTER_RENAME

Severity	Description
informational	Cluster with name '*old_name*' was renamed '*cluster.name*'.

HOST_DELETE

Severity	Description
informational	Host with name '*host.name*' was deleted.

CLUSTER_DELETE

Severity	Description
informational	Cluster with name '*cluster.name*' was deleted.

HOST_ADD_PORT

Severity	Description
informational	Port of type *type* and ID '*port_name*' was added to Host with name '*host.name*'.

CLUSTER_ADD_HOST

Severity	Description
informational	Host with name '*host.name*' was added to Cluster with name '*cluster.name*'.

HOST_REMOVE_PORT

Severity	Description
informational	Port of type *type* and ID '*port_name*' was removed from Host with name '*host.name*' was deleted.

CLUSTER_REMOVE_HOST

Severity	Description
informational	Host with name '*host.name*' was removed from Cluster with name '*cluster.name*'.

DESTINATION_DEFINE

Severity	Description
informational	Destination with name '*name*' was defined.

DESTINATION_UPDATE

Severity	Description
informational	Destination with name '*name*' was updated.

DESTINATION_DELETE

Severity	Description
informational	Destination with name '*name*' was deleted.

DESTINATION_RENAME

Severity	Description
informational	Destination with name '*old name*' was renamed '*new name*'.

DESTINATION_GROUP_CREATE

Severity	Description
informational	Destination Group with name '*name*' was created.

DESTINATION_GROUP_UPDATE

Severity	Description
informational	Destination Group with name '*name*' was updated.

DESTINATION_GROUP_DELETE

Severity	Description
informational	Destination Group with name '*name*' was deleted.

DESTINATION_GROUP_RENAME

Severity	Description
informational	Destination Group with name '*old name*' was renamed '*new name*'.

DESTINATION_GROUP_ADD_DESTINATION

Severity	Description
informational	Destination with name '*destination name*' was added to destination group '*destgroup name*'.

DESTINATION_GROUP_REMOVE_DESTINATION

Severity	Description
informational	Destination with name '*destination name*' was removed from destination group '*destgroup name*'.

RULE_CREATE

Severity	Description
informational	Rule with name '*name*' was created.

RULE_UPDATE

Severity	Description
informational	Rule with name '*name*' was updated.

RULE_DELETE

Severity	Description
informational	Rule with name '*name*' was deleted.

RULE_RENAME

Severity	Description
informational	Rule with name '*old name*' was renamed '*new name*'.

SMTP_GATEWAY_DEFINE

Severity	Description
informational	SMTP gateway with name '*name*' was defined.

SMTP_GATEWAY_UPDATE

Severity	Description
informational	SMTP gateway with name '*name*' was updated.

SMTP_GATEWAY_DELETE

Severity	Description
informational	SMTP gateway with name '*name*' was deleted.

SMTP_GATEWAY_RENAME

Severity	Description
informational	SMTP gateway with name '*old name*' was renamed '*new name*'.

SMTP_GATEWAY_PRIORITIZE

Severity	Description
informational	SMTP gateways were prioritized; the new order is *order*.

CALL_HOME_CONNECTION_OK

Severity	Description
informational	Events are sent to the Call Home server by SMTP gateway '*name*'.

CALL_HOME_CONNECTION_PROBLEM

Severity	Description	Troubleshooting
major	Events are not sent to the Call Home server by SMTP gateway '*name*'. Reason: *Event Reason*.	Please contact IBM support.

SMTP_GATEWAY_FAILED

Severity	Description
major	SMTP gateway with name '*name*' has failed. It will not be used until *Retry Time*.

SMTP_GATEWAY_VIA_NODE_FAILED

Severity	Description
warning	Sending event *Event Code* (*Event Index*) through *SMTP Gateway* via *Module ID* has failed; Error message: '*Error Message*'.

SMS_GATEWAY_DEFINE

Severity	Description
informational	SMS gateway with name '*name*' was defined.

SMS_GATEWAY_UPDATE

Severity	Description
informational	SMS gateway with name '*name*' was updated.

SMS_GATEWAY_DELETE

Severity	Description
informational	SMS gateway with name '*name*' was deleted.

SMS_GATEWAY_RENAME

Severity	Description
informational	SMS gateway with name '*old name*' was renamed '*new name*'.

SMS_GATEWAY_PRIORITIZE

Severity	Description
informational	SMS gateways were prioritized; the new order is *order*.

CONS_GROUP_CREATE

Severity	Description
informational	Consistency Group with name '*cg.name*' was created.

CONS_GROUP_CREATE_FAILED_TOO_MANY

Severity	Description	Troubleshooting
warning	Consistency Group with name '*name*' could not be created. You are attempting to add more Consistency Groups than the system permits.	Delete Consistency Groups to allow new ones to be created.

CONS_GROUP_RENAME

Severity	Description
informational	Consistency Group with name '*old_name*' was renamed '*cg.name*'.

SECONDARY_CONS_GROUP_RENAME

Severity	Description
informational	Consistency Group with name '*old_name*' was renamed '*cg.name*' by primary machine.

CONS_GROUP_DELETE

Severity	Description

Severity	Description
informational	Consistency Group with name '*cg.name*' was deleted.

CONS_GROUP_ADD_VOLUME

Severity	Description
informational	Volume with name '*volume.name*' was added to Consistency Group with name '*cg.name*'.

SLAVE_CONS_GROUP_ADD_VOLUME

Severity	Description
informational	Volume with name '*volume.name*' was added to Consistency Group with name '*cg.name*' by its remote peer.

CONS_GROUP_REMOVE_VOLUME

Severity	Description
informational	Volume with name '*volume.name*' was removed from Consistency Group with name '*cg.name*'.

SLAVE_CONS_GROUP_REMOVE_VOLUME

Severity	Description
informational	Volume with name '*volume.name*' was removed from Consistency Group with name '*cg.name*' by its remote peer.

CONS_GROUP_SNAPSHOTS_CREATE

Severity	Description
informational	Snapshot Group for Consistency Group with name '*cg.name*' was created with name '*sg.name*'.

CONS_GROUP_SNAPSHOTS_CREATE_FAILED_TOO_MANY

Severity	Description	Troubleshooting
warning	Snapshot Group for Consistency Group '*cg.name*' could not be created. You are attempting to add more Volumes than the system permits.	Delete Volumes to allow new ones to be created.

CONS_GROUP_SNAPSHOTS_OVERWRITE

Severity	Description
informational	Snapshot Group named '*sg.name*' was overriden for Consistency Group with name '*cg.name*'.

SLAVE_CONS_GROUP_SNAPSHOTS_CREATE

Severity	Description
informational	Mirrored Snapshot Group for Consistency Group with name '*cg.name*' was created with name '*sg.name*'.

HA_SLAVE_CONS_GROUP_SNAPSHOTS_CREATE

Severity	Description
informational	HyperSwap Snapshot Group for Consistency Group with name '*cg.name*' was created with name '*sg.name*'.

SLAVE_CONS_GROUP_SNAPSHOTS_OVERWRITE

Severity	Description
informational	Mirrored Snapshot Group named '*sg.name*' was overriden for Consistency Group with name '*cg.name*'.

HA_SLAVE_CONS_GROUP_SNAPSHOTS_OVERWRITE

Severity	Description
informational	HyperSwap Snapshot Group named '*sg.name*' was overriden for Consistency Group with name '*cg.name*'.

MIRROR_CONS_GROUP_SNAPSHOTS_CREATE

Severity	Description
informational	Mirrored Snapshot Group for Consistency Group with name '*cg.name*' was created with name '*sg.name*'.

HA_CONS_GROUP_SNAPSHOTS_CREATE

Severity	Description
informational	HyperSwap Snapshot Group for Consistency Group with name '*cg.name*' was created with name '*sg.name*'.

MIRROR_CONS_GROUP_SNAPSHOTS_OVERWRITE

Severity	Description
informational	Mirrored Snapshot Group named '*sg.name*' was overriden for Consistency Group with name '*cg.name*'.

HA_CONS_GROUP_SNAPSHOTS_OVERWRITE

Severity	Description
informational	HyperSwap Snapshot Group named '*sg.name*' was overriden for Consistency Group with name '*cg.name*'.

REMOTE_MIRROR_CONS_GROUP_SNAPSHOTS_NOT_CREATED_YET

Severity	Description
minor	Remote Mirrored Snapshot Group for Consistency Group with name '*remote CG name*' with name '*cg_sync_job.remote_snapgroup*' on Target '*target name*' were not created yet.

MIRROR_SNAPGROUP_CREATE_FAILED

Severity	Description
minor	Remote snapshot group named '*snapshot group name*' was not created successfully. Error code is '*error*'

SNAPSHOT_GROUP_RESTORE

Severity	Description
informational	Volumes were restored from Snapshot Group with name '*sg.name*'.

SNAPSHOT_GROUP_RENAME

Severity	Description
informational	Snapshot Group with name '*old_sg.name*' were renamed to '*sg.name*'.

SNAPSHOT_GROUP_DUPLICATE

Severity	Description
informational	All Snapshots in Snapshot Group with name '*sg.name*' were duplicated. Duplicate Snapshot Group is named '*sg.name*'.

SNAPSHOT_GROUP_FORMAT

Severity	Description
informational	All Snapshots in Snapshot Group with name '*sg.name*' were formatted'.

SNAPSHOT_GROUP_DELETE

Severity	Description
informational	All Snapshots in Snapshot Group with name '*sg.name*' were deleted.

SNAPSHOT_GROUP_CHANGE_PRIORITY

Severity	Description
informational	Deletion Priority of all Snapshots in Snapshot Group with name '*sg.name*' were changed from '*old priority*' to '*new priority*'.

SNAPSHOT_GROUP_LOCK

Severity	Description
informational	All Snapshots in Snapshot Group with name '*sg.name*' were locked.

SNAPSHOT_GROUP_UNLOCK

Severity	Description

Severity	Description
informational	All Snapshots in Snapshot Group with name '*sg.name*' were unlocked.

SNAPSHOT_GROUP_DELETED_DUE_TO_POOL_EXHAUSTION

Severity	Description
informational	All Snapshots in Snapshot Group with name '*snapshot.sg_name*' have been deleted because Storage Pool with name '*snapshot.pool_name*' is full.

SNAPSHOT_GROUP_DISBAND

Severity	Description
informational	Snapshot Group with name '*sg.name*' was dismantled. All Snapshots which belonged to that Snapshot Group should be accessed directly.

CONS_GROUP_MOVE

Severity	Description
informational	Consistency Group with name '*cg.name*' has been moved from Storage Pool '*orig_pool.name*' to Pool '*pool.name*'.

XCG_CREATE

Severity	Description
informational	Cross Consistency Group with name '*xcg*' was created.

XCG_DELETE

Severity	Description
informational	Cross Consistency Group with name '*xcg*' was deleted.

XCG_ADD_CG

Severity	Description
informational	CG with name '*cg.name*' was added to Cross Consistency Group with name '*xcg*'.

XCG_REMOVE_CG

Severity	Description
informational	CG with name '*cg.name*' was removed from Cross Consistency Group with name '*xcg*'.

TARGET_DEFINE

Severity	Description
informational	Target was defined named '*target.name*'.

TARGET_DEFINE_FAILED_TOO_MANY

Severity	Description	Troubleshooting
warning	Target could not be defined. You are attempting to define more targets than the system permits.	Delete targets to allow new ones to be defined.

TARGET_RENAME

Severity	Description
informational	Target named '*old_name*' was renamed '*target.name*'.

TARGET_DELETE

Severity	Description
informational	Target named '*target.name*' was deleted.

TARGET_ALLOW_ACCESS

Severity	Description
informational	Target '*target.name*' is allowed to access this machine.

TARGET_PORT_ADD

Severity	Description
informational	Port '*port_name*' was added to target named '*target.name*'.

TARGET_PORT_REMOVE

Severity	Description
informational	Port '*port_name*' was removed from target named '*target.name*'.

TARGET_PORT_ACTIVATE

Severity	Description
informational	Port '*port_name*' in target named '*target.name*' was activated.

TARGET_PORT_DEACTIVATE

Severity	Description
informational	Port '*port_name*' was deactivated in target named '*target.name*'.

TARGET_CONNECTIVITY_CREATE

Severity	Description
informational	Port '*Connection Remote Port Address*' of target named '*Connection Target Name*' is connected to the system through *Local FC Port*.

TARGET_ISCSI_CONNECTIVITY_CREATE

Severity	Description
informational	Port '*Connection Remote Port Address*' of target named '*Connection Target Name* is connected to the system through ip interface '*Local IP interface*'.

TARGET_CONNECTIVITY_CREATE_FAILED_TOO_MANY

Severity	Description	Troubleshooting
warning	Port could not be connected to the system. You are attempting to define more connections than the system permits.	Delete Connections to allow new ones to be created.

TARGET_CONNECTIVITY_DELETE

Severity	Description
informational	Port '*Connection Remote Port Address*' of target named '*Connection Target Name*' was disconnected from *Local FC Port*.

TARGET_ISCSI_CONNECTIVITY_DELETE

Severity	Description
informational	Port '*Connection Remote Port Address*' of target named '*Connection Target Name*' was disconnected from ip interface '*Local IP interface*'.

TARGET_CONNECTIVITY_ACTIVATE

Severity	Description
informational	Connectivity between Port '*Connection Remote Port Address*' of target named '*Connection Target Name*' and *Local FC Port* was activated.

TARGET_ISCSI_CONNECTIVITY_ACTIVATE

Severity	Description
informational	Connectivity between Port '*Connection Remote Port Address*' of target named '*Connection Target Name*' and ip interface '*Local IP interface*' was activated.

TARGET_CONNECTIVITY_DEACTIVATE

Severity	Description
informational	Connectivity between Port '*Connection Remote Port Address*' of target named '*Connection Target Name*' and *Local FC Port* was deactivated.

TARGET_ISCSI_CONNECTIVITY_DEACTIVATE

Severity	Description
informational	Connectivity between Port '*Connection Remote Port Address*' of target named '*Connection Target Name*' and ip interface '*Local IP interface*' was deactivated.

TARGET_CONNECTIVITY_CONFLICT_DETECTED

Severity	Description
major	Connectivity between Port '*Connection Remote Port Address*' of target named '*Connection Target Name*' and FC port '*Local IP interface*' will be deleted due to a connectivity conflict.

TARGET_ISCSI_CONNECTIVITY_CONFLICT_DETECTED

Severity	Description
major	Connectivity between Port '*Connection Remote Port Address*' of target named '*Connection Target Name*' and IP interface '*Local IP interface*' will be deleted due to a connectivity conflict.

TARGET_CONNECTION_ESTABLISHED

Severity	Description
informational	Target named '*target.name*' is accessible through remote service *module_id*.

TARGET_CONNECTION_DISCONNECTED

Severity	Description
minor	Target named '*target.name*' is no longer accessible through remote service *module_id*.

TARGET_DISCONNECTED

Severity	Description
major	Target named '*target.name*' is no longer accessible through any gateway module.

TARGET_CLOCK_SKEW_ABOVE_LIMIT

Severity	Description
warning	Target '*target.name*' has clock skew above the allowed limit relative to local machine.

TARGET_CLOCK_SKEW_RESOLVED

Severity	Description
informational	Target named '*target.name*' clock skew has been resolved.

TARGET_LINK_DOWN_BEYOND_THRESHOLD

Severity	Description
major	Target named '*target.name*' is not accessible for a long time.

OLVM_DELETE_ALL_REFERENCES_TO_SOURCE

Severity	Description
major	Target named '*target.name*' was released from all IBM Hyper-Scale Mobility relationships.

TARGET_SYNC_RATE_CHANGED

Severity	Description
informational	Target '*target.name*' sync rate changed. max_initialization_rate: '*target.max_initialization_rate*', max_resync_rate: '*target.max_resync_rate*', max_syncjob_rate: '*target.max_syncjob_rate*'.

TARGET_ADD_QUORUM_WITNESS

Severity	Description
Informational	Target '*target_name*' added Quorum Witness '*quorum_witness_name*'.

TARGET_REMOVE_QUORUM_WITNESS

Severity	Description
Informational	Target '*target_name*' removed Quorum Witness '*quorum_witness_name*'.

TARGET_SYSTEM_DETAILS_UPDATED

Severity	Description
Informational	Target named '*target_name*' has updated details: old system id '*old_system_id*', old machine serial '*old_machine_serial*', new system id '*new_system_id*', new machine serial '*new_machine_serial*'.

TARGET_HANDSHAKE_COMPLETED

Severity	Description
Informational	Target named '*target_name*' has completed handshake.

TARGET_HANDSHAKE_FAILED

Severity	Description
Major	Target named '*target_name*' handshake failed with reason '*Failure Reason*'

TARGET_HANDSHAKE_REINITIATED

Severity	Description
Informational	Target named '*target_name*' reinitiated handshake process.

SNAPSHOT_CREATE

Severity	Description
informational	Snapshot named '*snapshot.name*' was created for volume named '*volume.name*'.

SNAPSHOT_DELETE

Severity	Description
informational	Snapshot with name '*snapshot.name*' was deleted.

SNAPSHOT_OVERWRITE

Severity	Description
informational	Snapshot named '*snapshot.name*' was overriden for volume named '*volume.name*'.

SNAPSHOT_FORMAT

Severity	Description
informational	Snapshot named '*snapshot.name*' was formatted.

SNAPSHOT_CREATE_FAILED_TOO_MANY

Severity	Description	Troubleshooting
warning	Snapshot for volume named '*volume.name*' could not be created. You are attempting to add more volumes than the system permits.	Delete volumes to allow new ones to be created.

SNAPSHOT_DUPLICATE

Severity	Description
informational	Snapshot named '*snapshot.name*' was created as duplicate of Snapshot named '*original_snapshot.name*'.

SNAPSHOT_DUPLICATE_FAILED_TOO_MANY

Severity	Description	Troubleshooting
warning	Snapshot named '*snapshot.name*' could not be duplicated. You are attempting to add more volumes than the system permits.	Delete volumes to allow new ones to be created.

SNAPSHOT_RESTORE

Severity	Description

Severity	Description
informational	Volume named '*volume.name*' was restored from Snapshot named '*snapshot.name*'.

SNAPSHOT_CHANGE_PRIORITY

Severity	Description
informational	Snapshot Delete Priority of Snapshot named '*snapshot.name*' was changed from '*old_priority*' to '*snapshot.delete_priority*'.

SNAPSHOT_DELETED_DUE_TO_POOL_EXHAUSTION

Severity	Description
warning	Snapshot named '*snap.name*' has been deleted because Storage Pool named '*snap.pool_name*' is full.

MIRROR_SNAPSHOT_CREATE

Severity	Description
informational	Mirrored Snapshot named '*snapshot.name*' was created for volume named '*volume.name*'.

MIRROR_SNAPSHOT_CREATE_FAILED

Severity	Description
minor	Remote snapshot named '*snapshot name*' was not created successfully. Error code is '*error*'

MIRROR_SNAPSHOT_OVERWRITE

Severity	Description
informational	Mirrored Snapshot named '*snapshot.name*' was overriden for volume named '*volume.name*'.

MIRROR_SLAVE_SNAPSHOT_CREATE

Severity	Description
informational	Mirrored Snapshot named '*snapshot.name*' was created for volume named '*volume.name*'.

MIRROR_SLAVE_SNAPSHOT_OVERWRITE

Severity	Description
informational	Mirrored Snapshot named '*snapshot.name*' was overriden for volume named '*volume.name*'.

MEDIUM_ERROR_IN_DATA_MIGRATION

Severity	Description	Troubleshooting
critical	Medium error in data migration into volume '*Volume Name*' at LBA *LBA* for *Length* blocks.	Remote machine indicated Medium Error when read.

TRANSACTION_NODE_DOES_NOT_USE_OPTIMAL_SRP_PATH

Severity	Description	Troubleshooting
warning	Data service '*service*' does not use the optimal path to '*enclosure*'.	Contact IBM Support

TRANSACTION_NODE_USES_OPTIMAL_SRP_PATH

Severity	Description	Troubleshooting
informational	Data service '*service*' uses the optimal path to '*enclosure*'.	Contact IBM Support

USER_DEFINED

Severity	Description
informational	A user with name '*Name*' and category *Category* was defined.

USER_DELETED

Severity	Description
informational	A user with name '*Name*' and category *Category* was deleted.

USER_RENAMED

Severity	Description
informational	User with name '*Old Name*' was renamed '*New Name*'.

USER_UPDATED

Severity	Description
informational	User with name '*Name*' was updated.

USER_ADDED_TO_USER_GROUP

Severity	Description
informational	User '*User Name*' was added to user group '*User Group Name*'.

USER_REMOVED_FROM_USER_GROUP

Severity	Description
informational	User '*User Name*' was removed from user group '*User Group Name*'.

USER_GROUP_CREATED

Severity	Description
informational	A user group with name '*Name*' was created.

USER_GROUP_DELETED

Severity	Description
informational	A user group with name '*Name*' was deleted.

USER_GROUP_RENAMED

Severity	Description
informational	User group with name '*Old Name*' was renamed '*New Name*'.

LDAP_AUTHENTICATION_ACTIVATED

Severity	Description
informational	LDAP authentication activated.

LDAP_AUTHENTICATION_DEACTIVATED

Severity	Description
warning	LDAP authentication deactivated.

LDAP_CONFIGURATION_CHANGED

Severity	Description
warning	LDAP configuration has changed.

LDAP_CONFIGURATION_RESET

Severity	Description
warning	LDAP configuration has reset.

USER_LOGIN_HAS_SUCCEEDED

Severity	Description
informational	User '*User Name*' from IP '*Client Address*' successfully logged into the system.

USER_LOGIN_HAS_FAILED

Severity	Description
warning	User '*User Name*' from IP '*Client Address*' failed logging into the system.

USER_HAS_FAILED_TO_RUN_COMMAND

Severity	Description
warning	User '*User Name*' from IP '*Client Address*' failed authentication when trying to run command '*Command Line*'.

LDAP_SERVER_INACCESSIBLE

Severity	Description
minor	LDAP server *FQDN* is inaccessible.

LDAP_SERVER_ACCESSIBLE

Severity	Description
informational	LDAP server *FQDN* is now accessible.

LDAP_SSL_CERTIFICATE_ABOUT_TO_EXPIRE

Severity	Description
variable	SSL Certificate of LDAP server '*Server FQDN*' is about to expire on *Expiration Date* (*Counter* notification).

LDAP_SERVER_WAS_ADDED

Severity	Description
informational	LDAP server '*Server FQDN*' was added to the system.

LDAP_SERVER_WAS_REMOVED

Severity	Description
informational	LDAP server '*Server FQDN*' was removed from the system.

DESIGNATED_MSM_USER

Severity	Description
informational	*Description*

DOMAIN_POLICY_SET

Severity	Description
informational	Domain policy for *Parameter Name* set to '*Parameter Value*'

USER_ADDED_TO_DOMAIN

Severity	Description
informational	User *User Name* was added to domain *Domain Name* (*Exclusive*).

USER_REMOVED_FROM_DOMAIN

Severity	Description
informational	User *User Name* was removed from domain *Domain Name*.

APPADMIN_CAPABILITIES_SET

Severity	Description
informational	Application admin capabilities have been set to *Capabilities*

ACCESS_TO_HOST_GRANTED_TO_USER_GROUP

Severity	Description
informational	User group '*User Group Name*' was granted access to host '*Host Name*'.

ACCESS_OF_USER_GROUP_TO_HOST_REMOVED

Severity	Description
informational	Access of User group '*User Group Name*' to host '*Host Name*' was removed.

ACCESS_TO_CLUSTER_GRANTED_TO_USER_GROUP

Severity	Description
informational	User group '*User Group Name*' was granted access to cluster '*Cluster Name*'.

ACCESS_OF_USER_GROUP_TO_CLUSTER_REMOVED

Severity	Description
informational	Access of User group '*User Group Name*' to cluster '*Cluster Name*' was removed.

COMPONENT_TEST_HAS_FAILED

Severity	Description	Troubleshooting
variable	Test of *Component ID* has failed. Failure reason: *Failure Reason*.	Contact IBM Support

COMPONENT_TEST_SUCCEEDED

Severity	Description	Troubleshooting
informational	Test of *Component ID* succeeded.	Contact IBM Support

MODULE_COMPONENT_TEST_STARTED

Severity	Description
informational	Test of *Component ID* started.

DISK_COMPONENT_TEST_STARTED

Severity	Description
informational	Test of *Component ID* started.

IB_SWITCH_COMPONENT_TEST_STARTED

Severity	Description
informational	Test of *Component ID* started.

SSD_COMPONENT_TEST_STARTED

Severity	Description
informational	Test of *Component ID* started.

VAULT_DEVICE_COMPONENT_TEST_STARTED

Severity	Description
informational	Test of *Component ID* started.

BOOT_MEDIA_COMPONENT_TEST_STARTED

Severity	Description
informational	Test of *Component ID* started.

FLASH_COMPONENT_TEST_STARTED

Severity	Description
informational	Test of *Component ID* started.

BOOT_MEDIA_FAILED

Severity	Description	Troubleshooting
major	*Component ID* has failed. Hardware status: *Status*.	Contact IBM Support

MODULE_BBU_TEST_STARTED

Severity	Description
informational	Test of *Component ID* started.

MODULE_BBU_TEST_WILL_CONTINUE

Severity	Description
warning	Test of *Component ID* will continue. After waiting *Minutes* minutes, capacity is still *Capacity*%.

COMPONENT_WAS_PHASED_OUT

Severity	Description
informational	*Component ID* was phased-out.

COMPONENT_WAS_FAILED

Severity	Description
variable	Component *Component ID* was marked as failed.

COMPONENT_FAILURE_WAS_CANCELED

Severity	Description
informational	Component *Component ID* failure status was reset.

COMPONENT_WAS_PHASED_IN

Severity	Description
informational	*Component ID* was phased-in.

COMPONENT_WAS_EQUIPPED

Severity	Description
informational	*Component ID* was equipped.

INTERFACE_SERVICES_ACTIVATED

Severity	Description
informational	Interface services of *Module ID* were activated.

COMPONENT_FIRMWARE_UPGRADE_ABORTING

Severity	Description
warning	Aborting *Upgrade type* upgrade of *Firmware type* firmware, version *Label*, on *Scope*. Abort reason: *Reason*. Waiting for current upgrade item to complete.

COMPONENT_FIRMWARE_UPGRADE_ABORTED

Severity	Description
warning	Aborted *Upgrade type* upgrade of *Firmware type* firmware, version *Label*, on *Scope*. Abort reason: *Reason*. Progress *Attempted*/*Total*, *Successes* succeeded, *Failures* failed, *No-Ops* no-ops.

COMPONENT_FIRMWARE_UPGRADE_DONE

Severity	Description
informational	Finished *Upgrade type* upgrade of *Firmware type* firmware, version *Label*, on *Scope*. *Successes* succeeded, *Failures* failed, *No-Ops* no-ops.

COMPONENT_FIRMWARE_UPGRADE_STARTED

Severity	Description
informational	Starting *Upgrade type* upgrade of *Firmware type* firmware, version *Label*, on *Scope*.

COMPONENT_FIRMWARE_CANNOT_PHASEOUT_COMPONENT

Severity	Description
minor	Cannot phase out *Component ID*: *Error*. Firmware upgrade result was: *Upgrade result*.

COMPONENT_FIRMWARE_CANNOT_FAIL_COMPONENT

Severity	Description
minor	Cannot fail *Component ID*: *Error*. Firmware upgrade result was: *Upgrade result*.

MIRRORING_CONNECTIVITY_TO_NON_XIV_TARGET

Severity	Description	Troubleshooting
warning	Gateway Node #*Node ID*: connection to *target name*:*target's connection index* mirroring connection was established, but being ignored because the remote end is not an XIV target or is not properly configured	Please make sure the target's designation is correct, that the connection's parameters identify the intended system and that the intended system has a target_port defined for this system.

DM_CONNECTIVITY_TO_XIV_TARGET

Severity	Description	Troubleshooting
warning	Gateway Node #*Node ID*: connection to *target name*:*target's connection index* DM connection was established, but being ignored because the remote end is an XIV target configured for mirroring, rather than a host	Please make sure the target's designation is correct, that the connection's parameters identify the intended system and that the intended system has a host defined for this system (and not a target_port).

EMERGENCY_ROOT_ACCESS

Severity	Description
warning	Emergency login to 'root' account on module '*Component ID*' from '*IP Address*' using key number '*Authorized Key Number*'.

EMERGENCY_CONSOLE_ACCESS

Severity	Description
warning	Emergency login to '*Unix Account Name*' account on module '*Component ID*' from tty '*TTY Device*'.

CR_BYPASS_ACCESS

Severity	Description
warning	*Command that bypasses CR mechanism* access to '*Unix Account Name*' account on module '*Component ID*' from '*IP Address*'.

CR_KEY_SETUP_OK

Severity	Description
informational	Challenge-response key was successfully set on all modules in the system.

CR_KEY_UPGRADE_NOT_DONE

Severity	Description
warning	Challenge-response key was not upgraded on the system since a valid key has been previously set.

CR_KEY_SETUP_FAILED

Severity	Description
major	Failed to set challenge-response key on module '*Component ID*'.

SSH_REVOKE_KEY_OK

Severity	Description
informational	Authorized SSH key ending with '*Tail of Authorized SSH key*' was successfully revoked for user '*Unix Account Name*' on all modules in the system.

SSH_REVOKE_KEY_FAILED

Severity	Description
major	Failed to revoke authorized SSH key ending with '*Tail of Authorized SSH key*' for user '*Unix Account Name*' on module '*Component ID*'.

IB_SWITCH_PHASEOUT_STARTED

Severity	Description
informational	System started phasing out *Component ID*.

IB_SWITCH_PHASEIN_STARTED

Severity	Description
informational	System started phasing in *Component ID*.

IB_SWITCH_PHASEIN_FAILED

Severity	Description
warning	*Component ID* has failed to phase-in.

IB_SWITCH_CONFIG_FAILED

Severity	Description	Troubleshooting
warning	*Component ID* could not be configured	Contact IBM Support

IB_SWITCH_FIRMWARE_INCOMPATIBLE

Severity	Description	Troubleshooting
warning	The firmware version of *Component ID* is '*New Version*'. It should be '*Old Version*'.	None

IB_SWITCH_CPLD_INCOMPATIBLE

Severity	Description	Troubleshooting
warning	The CPLD version of *Component ID* is invalid.	None

IB_SWITCH_FIRMWARE_UPDATE_IN_PROGRESS

Severity	Description	Troubleshooting
informational	Firmware version of *Component ID* is '*Old Version*'. It should be '*New Version*'. Firmware will be updated. It may take a while.	Wait for IB switch to complete initialization.

IB_SWITCH_FIRMWARE_UPDATED

Severity	Description	Troubleshooting
informational	The firmware version of *Component ID* was updated to '*New Version*'.	None.

IB_SWITCH_LOG_COLLECT_OK

Severity	Description
informational	Log collection for IB switch '*switch_id*' completed successfuly. Log can be found in module '*log_module*' in the following directory: '*log_location*'.

IB_SWITCH_LOG_COLLECT_FAILED

Severity	Description
warning	Log collection for IB switch '*switch_id*' failed. Failure reason: '*failure_reason*'.

IB_SWITCH_MGMT_LINK_AVAILABLE

Severity	Description
informational	Management link *Type* of *Component ID* is available.

IB_SWITCH_MGMT_LINK_UNAVAIL

Severity	Description
warning	Management link *Type* of *Component ID* is unavailable.

IB_SWITCH_MGMT_LINK_MISWIRED

Severity	Description	Troubleshooting
warning	Management link *Type* of *Component ID* is connected to *Router* instead of *Expected Router*.	Check wiring

IB_SWITCH_MGMT_AVAILABLE

Severity	Description
informational	Management of *Component ID* is available.

IB_SWITCH_MGMT_UNAVAIL

Severity	Description
major	Management of *Component ID* is unavailable.

IB_SWITCH_PSU_OK

Severity	Description
informational	*Component ID* has returned to normal state.

IB_SWITCH_PSU_MONITOR_FAILED

Severity	Description
minor	*Component ID* sensor cannot be read.

IB_SWITCH_PSU_IS_MISSING

Severity	Description
minor	*Component ID* is not present.

IB_SWITCH_PSU_FAIL

Severity	Description
major	*Component ID* failed.

IB_SWITCH_BBU_OK

Severity	Description
informational	*Component ID* has returned to normal state.

IB_SWITCH_BBU_MONITOR_FAILED

Severity	Description
minor	*Component ID* sensor cannot be read.

IB_SWITCH_BBU_IS_MISSING

Severity	Description
minor	*Component ID* is not present.

IB_SWITCH_BBU_FAIL

Severity	Description
major	*Component ID* failed.

IB_SWITCH_BBU_END_OF_LIFE

Severity	Description
major	*Component ID* has reached end of life.

IB_SWITCH_FAN_OK

Severity	Description
informational	*Component ID* has returned to normal state.

IB_SWITCH_FAN_MONITOR_FAILED

Severity	Description
minor	*Component ID* sensor cannot be read.

IB_SWITCH_FAN_IS_MISSING

Severity	Description
minor	*Component ID* is not present.

IB_SWITCH_FAN_FAIL

Severity	Description
major	*Component ID* failed.

IB_SWITCH_PSU_FAN_OK

Severity	Description
informational	*Component ID* has returned to normal state.

IB_SWITCH_PSU_FAN_MONITOR_FAILED

Severity	Description
minor	*Component ID* sensor cannot be read.

IB_SWITCH_PSU_FAN_IS_MISSING

Severity	Description
minor	*Component ID* is not present.

IB_SWITCH_PSU_FAN_FAIL

Severity	Description
major	*Component ID* failed.

IB_SWITCH_VOLTAGE_MONITOR_FAILED

Severity	Description
minor	*Component ID* voltage sensor *Sensor Type* cannot be read.

IB_SWITCH_VOLTAGE_CHANGE

Severity	Description
informational	*Component ID* voltage sensor *Sensor Type* changed from *Old Status* to *Status*.

IB_SWITCH_PSU_VOLTAGE_MONITOR_FAILED

Severity	Description
minor	*Component ID* voltage sensor cannot be read.

IB_SWITCH_PSU_VOLTAGE_CHANGE

Severity	Description
informational	*Component ID* voltage sensor changed from *Old Status* to *Status*.

IB_SWITCH_BBU_VOLTAGE_MONITOR_FAILED

Severity	Description
minor	*Component ID* voltage sensor cannot be read.

IB_SWITCH_BBU_VOLTAGE_CHANGE

Severity	Description
informational	*Component ID* voltage sensor changed from *Old Status* to *Status*.

IB_SWITCH_TEMPERATURE_MONITOR_FAILED

Severity	Description
minor	*Component ID* sensor *Sensor Type* temperature cannot be read.

IB_SWITCH_TEMPERATURE_OK

Severity	Description
informational	*Component ID* sensor *Sensor Type* temperature has returned to normal state.

IB_SWITCH_TEMPERATURE_HIGH

Severity	Description
minor	*Component ID* sensor *Sensor Type* temperature is high.

IB_SWITCH_TEMPERATURE_CRITICAL

Severity	Description
major	*Component ID* sensor *Sensor Type* temperature is critical.

IB_SWITCH_PSU_TEMPERATURE_MONITOR_FAILED

Severity	Description
minor	*Component ID* temperature cannot be read.

IB_SWITCH_PSU_TEMPERATURE_OK

Severity	Description
informational	*Component ID* temperature has returned to normal state.

IB_SWITCH_PSU_TEMPERATURE_HIGH

Severity	Description
minor	*Component ID* temperature is high.

IB_SWITCH_PSU_TEMPERATURE_CRITICAL

Severity	Description
major	*Component ID* temperature is critical.

IB_SWITCH_BBU_TEMPERATURE_MONITOR_FAILED

Severity	Description
minor	*Component ID* temperature cannot be read.

IB_SWITCH_BBU_TEMPERATURE_OK

Severity	Description
informational	*Component ID* temperature has returned to normal state.

IB_SWITCH_BBU_TEMPERATURE_HIGH

Severity	Description
minor	*Component ID* temperature is high.

IB_SWITCH_BBU_TEMPERATURE_CRITICAL

Severity	Description
major	*Component ID* temperature is critical.

IB_PORT_MOVED

Severity	Description
informational	Infinibind module port '*module_port*' moved from '*from_port*' to '*to_port*'.

SYSTEM_TEMPERATURE_IS_ABOVE_NORMAL

Severity	Description	Troubleshooting
warning	System temperature is *System Temperature*C, which is above the normal temperature.	Cool the system down.

SYSTEM_TEMPERATURE_IS_HIGH

Severity	Description	Troubleshooting
minor	System temperature is *System Temperature*C, which is high.	Cool the system down.

SYSTEM_TEMPERATURE_IS_CRITICALLY_HIGH

Severity	Description	Troubleshooting
critical	System temperature is *System Temperature*C, which exceeds operational level. Please initiate shutdown sequence. Without further action, the system will automatically shut itself down if it reaches *Shutdown Threshold*C.	Cool the system down immediately or shut down the system using 'shutdown -y' and contact IBM support.

SYSTEM_TEMPERATURE_IS_CRITICALLY_HIGH_SHUTTING_DOWN

Severity	Description	Troubleshooting
critical	System temperature is *System Temperature*C, which is critically high. Shutting down the system.	Shut down the system using 'shutdown -y' and contact IBM support.

SYSTEM_TEMPERATURE_IS_TOO_HIGH

Severity	Description	Troubleshooting
major	System temperature is *System Temperature*C. It approaches the maximal allowable value.	Cool the system down and contact IBM support.

SYSTEM_TEMPERATURE_IS_TOO_LOW

Severity	Description	Troubleshooting
major	System temperature is *System Temperature*C, which is lower than the minimal allowable value.	Contact IBM Support

SYSTEM_TEMPERATURE_IS_OK_NOW

Severity	Description
informational	System temperature is *System Temperature*C, which is within allowed limits.

SYSTEM_AVERAGE_POWER_PREPARATION_STARTED

Severity	Description
informational	System average power consumption preparation has started, when it's over the up-to-date value will be available.

SYSTEM_AVERAGE_POWER_PREPARATION_OVER

Severity	Description
informational	System average power consumption preparation is over, you can now read the up-to-date value.

SYSTEM_AVERAGE_POWER_PREPARATION_FAILED

Severity	Description
warning	System average power consumption preparation has failed, try again later.

ENCRYPT_ENABLE_DRIVE_FAILED

Severity	Description	Troubleshooting
major	Failed to enable encryption for *Component ID*. Error code: *Failure Reason*.	Contact IBM Support

ENCRYPT_ENABLE_VAULT_DEVICE_FAILED

Severity	Description	Troubleshooting
major	Failed to enable encryption for *Component ID*. Error code: *Failure Reason*.	Contact IBM Support

VAULT_DEVICE_ENCRYPTING_ENABLE_FAILED

Severity	Description	Troubleshooting
major	Failed to enable encryption for *Component ID*. Error code: *Failure Reason*.	Contact IBM Support

VAULT_DEVICE_SECURE_ERASE_FAILED

Severity	Description	Troubleshooting
major	Failed to secure erase *Component ID*. Error code: *Failure Reason*.	Contact IBM Support

VAULT_DEVICE_SECURE_ERASE_SUCCESSFUL

Severity	Description	Troubleshooting
informational	Secure erase was successful for *Component ID*.	Contact IBM Support

VAULT_DEVICE_SECURE_ERASE_NOT_DONE

Severity	Description	Troubleshooting
informational	Secure erase was not done for *Component ID* as it is unsecured.	Contact IBM Support

MODULE_BBU_OVERHEATING

Severity	Description	Troubleshooting
major	*bbu* reached a temperature of *temp*C, above *critical_set*C. BBU will be disconnected until it cools down below *critical_release*C.	Contact IBM Support

MODULE_BBU_IS_DISCHARGING

Severity	Description
informational	*BBU id* changed state from '*old_state*' to '*new state*'.

MODULE_BBU_STOPPED_DISCHARGING

Severity	Description
informational	*BBU id* changed state from '*old_state*' to '*new state*'.

MODULE_BBU_CHARGING_WAS_EXPLICITLY_ENABLED

Severity	Description
minor	*BBU id* was not charging, it had to be reset explicitly.

MODULE_BBU_NOT_CHARGING_AFTER_RESET

Severity	Description
major	*BBU id* is still not charging after *Reset Attempts* reset attempts.

MODULE_BBU_NOT_CHARGED_AFTER_RESET

Severity	Description
major	*BBU id* was reset and it is charging, but after *Minutes* minutes it is still only *Percent Charged*% charged, which is not enough.

MODULE_BBU_STILL_NOT_SUFFICIENTLY_CHARGED

Severity	Description
major	*BBU id* is charging, but after *Minutes* minutes it is still only *Percent Charged*% charged, which is not enough.

MODULE_BBU_DISCHARGING_WAS_EXPLICITLY_ENABLED

Severity	Description
minor	*BBU id* was in discharge disabled mode, it had to be enabled explicitly.

MODULE_BBU_IS_CHARGING

Severity	Description
informational	*BBU id* changed state from '*old_state*' to '*new state*'.

MODULE_BBU_IS_FULL

Severity	Description
informational	*BBU id* changed state from '*old_state*' to '*new state*'.

MODULE_BBU_DRIVER_NOT_LOADED

Severity	Description	Troubleshooting
major	BBU driver is not loaded in module *Component*. modprobe ruby	Contact IBM Support

MODULE_BBU_CONTROLLER_NOT_PRESENT

Severity	Description	Troubleshooting
major	BBU controller board not detected in module *Component*.	Contact IBM Support

MODULE_IS_MISSING_EPOW_CABLE

Severity	Description	Troubleshooting
major	The EPOW cable in the BBU controller board on module *Component* is not detected.	Contact IBM Support

MODULE_EPOW_CABLE_OK_NOW

Severity	Description
informational	The EPOW cable in the BBU controller board on module *Component* is now OK.

MODULE_IS_MISSING_POWER_SENSE_CABLE

Severity	Description	Troubleshooting
major	The power sense cable in the BBU controller board on module *Component* is not detected.	Contact IBM Support

MODULE_POWER_SENSE_CABLE_OK_NOW

Severity	Description
informational	The power sense cable in the BBU controller board on module *Component* is now OK.

MODULE_BBU_CALIBRATION_STARTED

Severity	Description
informational	*BBU id* started calibration.

MODULE_BBU_CALIBRATION_ENDED

Severity	Description
informational	*BBU id* ended calibration with status '*result*'.

MODULE_BBU_TEMPERATURE_TOO_HIGH_FOR_CALIBRATION

Severity	Description	Troubleshooting
major	*BBU id* temperature is *Temperature.Temperature Tenths*C which is too high. Calibration is stopped.	Contact IBM Support

MODULE_BBU_TEST_IN_CHARGING_MODE

Severity	Description
informational	Test of *Component ID* is pending and will resume once it gets to a capacity of *Target Capacity*%. The current capacity is *Current Capacity*%.

SDR_PSU_STATUS_OK

Severity	Description
informational	*Psu* (location *Location*) is now OK. Changed from '*previous_sdr_status*' to '*sdr_status*'.

SDR_PSU_STATUS_BAD

Severity	Description
warning	*PSU* (location *Location*) is failed or off. Changed from '*previous_sdr_status*' to '*current_sdr_status*'.

INVALID_PSU_PART_NUMBER

Severity	Description	Troubleshooting
major	*PSU* has an invalid part number '*PN*'.	Please contact IBM support and have the PSU replaced.

CMOS_BATTERY_TOO_WEAK

Severity	Description	Troubleshooting
major	The CMOS battery on *Module* is too weak.	Please contact IBM support and have the battery replaced.

CMOS_BATTERY_IS_OK

Severity	Description
informational	The CMOS battery on *Module* is now OK.

FC_LINK_IS_NOW_DOWN

Severity	Description	Troubleshooting
major	FC port *Component* Active Firmware *Firmware version* - link disconnected.	Contact IBM Support

FC_LINK_IS_NOW_UP

Severity	Description	Troubleshooting
informational	FC port *Component* - link regained.	Contact IBM Support

FC_LINK_SYNC_ERROR

Severity	Description	Troubleshooting
major	FC port *Component* - errors on the physical layer: *Reason*.	Please contact support.

FC_PORT_TEST_STARTED

Severity	Description
informational	FC port *Component* - test started

FC_PORT_TEST_NOT_STARTED

Severity	Description	Troubleshooting
informational	FC port *Component* - test not started.	Check port state

FC_PORT_TEST_FAILED

Severity	Description	Troubleshooting
major	FC port *Component* - test failed.	Contact IBM Support

FC_PORT_TEST_SUCCESS

Severity	Description
informational	FC port *Component* - test success.

FC_PORT_TEST_ABORTED

Severity	Description
informational	FC port *Component* - test aborted.

COMPONENT_NETWORK_LINK_IS_DOWN

Severity	Description	Troubleshooting
major	Network interface to *Connected Component* on *Component ID* - link disconnected.	Contact IBM Support

COMPONENT_NETWORK_LINK_IS_UP

Severity	Description	Troubleshooting
informational	Network interface to component *Connected Component* on *Component ID* - link regained.	Contact IBM Support

MODULE_IS_MISSING_REQUIRED_MEMORY

Severity	Description	Troubleshooting
major	*Component ID* has less memory (*actual_mem* GB) than is defined for use (*req_mem* GB).	Please contact your Administrator.

POD_IB_MISWIRE

Severity	Description	Troubleshooting
warning	POD module miswired: *Module ID*.	Contact IBM Support

POD_IB_MISWIRE_CORRECTED

Severity	Description
informational	POD module miswire corrected: *Module ID*.

MODULE_SET_LED_LOCATOR_FAILED

Severity	Description
warning	Failed to set LED locator on *module*.

MODULE_SET_LED_LOCATOR_COMPLETED

Severity	Description
informational	LED locator set successfully on *module*.

PERF_CLASS_RESOURCE_EXHAUSTION

Severity	Description
warning	Exhausted all allowed resources for performance classes on *Module Id*, BUSY until resources available.

CONNECTED_HOSTS_LIMIT_REACHED

Severity	Description
informational	Number of connected Hosts was reached for port '*port_id*' in Module *Module Id*.

QoS_HAS_BEEN_TRIGGERED

Severity	Description
informational	Queues on port '*port_id*' in Module *Module Id* caused QoS to be activated.

PERF_CLASS_RATE_AT_LIMIT

Severity	Description
informational	Performance class '*perf_class*' on *Module Id* reached its limit of *Limit Limit Name*, IOs being throttled.

INDEPENDENT_PERF_CLASS_RATE_AT_LIMIT

Severity	Description
informational	Performance class '*perf_class*' object *type:name* on *Module Id* reached its limit of *Limit Limit Name*, IOs being throttled.

PORT_PREP_FOR_UPGRADE_TIMED_OUT

Severity	Description
warning	Preparation of *port_type* port '*local_port_name*' for hot-upgrade timed out due to host '*host_name*' port '*host_port_name*'*host_port_addr*

INTERFACE_DISCONNECTED_FROM_TARGET

Severity	Description
major	Interface node on module *module* cannot access target '*target*' through any gateway module.

INTERFACE_RECONNECTED_TO_TARGET

Severity	Description
major	Interface node on module *module* can access target '*target*'.

METADATA_SERVICE_DB_CREATE

Severity	Description
informational	Database *DB* was created

METADATA_SERVICE_DB_DELETE

Severity	Description
informational	Database *DB* was deleted

IPINTERFACE_CREATE

Severity	Description
informational	A new iSCSI IP interface was defined with name '*interface name*' on module *module* with port '*port list*' and IP address *IP address*.

IPINTERFACE_DELETE

Severity	Description
informational	iSCSI IP interface with name '*interface name*' was deleted.

IPINTERFACE_RENAME

Severity	Description
informational	iSCSI IP interface with name '*old name*' and was renamed '*interface name*'.

IPINTERFACE_UPDATE

Severity	Description
informational	ISCSI IP interface with name '*interface name*' was updated. Its IP address is *IP address*.

IPINTERFACE_UPDATE_MANAGEMENT

Severity	Description
informational	Management IP interfaces were updated. Management IPs are *IP addresses*.

IPINTERFACE_UPDATE_MANAGEMENT_IPV6

Severity	Description
informational	Management IP interfaces were updated. Management IPv6 addresses are *IPv6 addresses*.

IPINTERFACE_UPDATE_VPN

Severity	Description
informational	VPN IP interfaces were updated. VPN IPs are *IP addresses*.

IPINTERFACE_VLAN_UPDATE

Severity	Description
informational	System VLAN configuration was updated: PCP=*PCP*

IPINTERFACE_UPDATE_VPN_IPV6

Severity	Description
informational	VPN IPv6 interfaces were updated. VPN IPv6 addresses are *IP addresses*.

ETHERNET_PORT_UPDATE

Severity	Description
informational	Ethernet port with name *Ethernet port* was updated.

AUXILIARY_INTERNAL_PORTS_ENABLED

Severity	Description
informational	*Port Count* auxiliary internal Ethernet ports were enabled.

AUXILIARY_INTERNAL_PORTS_DISABLED

Severity	Description
informational	*Port Count* auxiliary internal Ethernet ports were disabled.

IPSEC_ENABLED

Severity	Description
informational	IPSec was enabled

IPSEC_DISABLED

Severity	Description
informational	IPSec was disabled

IPSEC_CONNECTION_ADDED

Severity	Description
informational	A new IPSec connection named '*name*' was added

IPSEC_CONNECTION_UPDATED

Severity	Description
informational	The IPSec connection named '*name*' was updated

IPSEC_CONNECTION_REMOVED

Severity	Description
informational	The IPSec connection named '*name*' was removed

PRIVATE_KEY_ADDED

Severity	Description
informational	A new private key named '*name*' with fingerprint '*fingerprint*' and size *key_size* bits was added.

CERTIFICATE_REMOVED

Severity	Description
informational	The certificate named '*name*' was removed.

PKCS12_CERTIFICATE_ADDED

Severity	Description
informational	A new PKCS#12 named '*name*' with fingerprint '*fingerprint*' was added.

PKI_RENAME

Severity	Description
informational	PKI with the name '*old name*' was renamed to '*new name*'

PKI_UPDATED

Severity	Description
informational	PKI with the name '*name*' and fingerprint '*fingerprint*' was updated

PROTOCOL_CONFIGURATION_CHANGED

Severity	Description
informational	Protocol configuration changed for protocol '*Protocol Type*'.

EMAIL_HAS_FAILED

Severity	Description	Troubleshooting
variable	Sending event *Event Code* (*Event Index*) to *Destination List* via *SMTP Gateway* failed. Module: *Module ID*; Error message: '*Error Message*'; timeout expired: *Timeout Expired?*.	Contact IBM Support

BULK_EMAIL_HAS_FAILED

Severity	Description	Troubleshooting
variable	Sending bulk email with *Events Number* events to *Destination List* via *SMTP Gateway* failed. Module: *Module ID*; Error message: '*Error Message*'; timeout expired: *Timeout Expired?*.	Contact IBM Support

SMS_HAS_FAILED

Severity	Description	Troubleshooting
variable	Sending event *Event Code* (*Event Index*) to *Destination List* via *SMS Gateway* and *SMTP Gateway* failed. Module: *Module ID*; Error message: '*Error Message*'; timeout expired: *Timeout Expired?*.	Contact IBM Support

HTTPS_HAS_FAILED

Severity	Description	Troubleshooting
variable	Sending event *Event Code* (*Event Index*) to *Destination List* via *HTTPS address* failed. Module: *Module ID*; Error message: '*Error Message*' (*HTTP error code*); timeout expired: *Timeout Expired?*.	Contact IBM Support

EMAIL_NOT_SENT

Severity	Description	Troubleshooting
variable	Sending event *Event Code* (*Event Index*) to *Destination List* via *SMTP Gateway* was waived because of failed SMTP gateway. It will be not be used until *Retry Time*.	Contact IBM Support

SMS_NOT_SENT

Severity	Description	Troubleshooting
variable	Sending event *Event Code* (*Event Index*) to *Destination List* via *SMS Gateway* and *SMTP Gateway* was waived because of failed SMTP gateway. It will be not be used until *Retry Time*.	Contact IBM Support

HEARTBEAT_EMAIL_HAS_FAILED

Severity	Description	Troubleshooting
minor	Sending heartbeat to *Destination Name* via *SMTP Gateway* failed. Module: *Module ID*; Error message: '*Error Message*'; timeout expired: *Timeout Expired?*.	Contact IBM Support

HEARTBEAT_SMS_HAS_FAILED

Severity	Description	Troubleshooting
minor	Sending heartbeat to *Destination Name* via *SMS Gateway* and *SMTP Gateway* failed. Module: *Module ID*; Error message: '*Error Message*'; timeout expired: *Timeout Expired?*.	Contact IBM Support

TEST_EMAIL_HAS_FAILED

Severity	Description	Troubleshooting
minor	Sending test to *Destination Name* via *SMTP Gateway* failed. Module: *Module ID*; Error message: '*Error Message*'; timeout expired: *Timeout Expired?*.	Contact IBM Support

TEST_SMS_HAS_FAILED

Severity	Description	Troubleshooting
minor	Sending test to *Destination Name* via *SMS Gateway* and *SMTP Gateway* failed. Module: *Module ID*; Error message: '*Error Message*'; timeout expired: *Timeout Expired?*.	Contact IBM Support

CUSTOM_EVENT

Severity	Description
variable	*Description*

UPGRADE_SOFTWARE_DOWNLOAD_FINISHED

Severity	Description
informational	Finished downloading software needed for upgrade to version *version*. Upgrade consequence is *consequence*

UPGRADE_FILE_LIST_RETRIEVAL_FAILED

Severity	Description	Troubleshooting
critical	Could not receive new version's file list from repository. Error code is *error*.	Contact IBM Support

UPGRADE_STARTS

Severity	Description
informational	System starting an upgrade.

PRE_UPGRADE

Severity	Description
informational	System preparing an upgrade procedure type *type* .

UPGRADE_IS_OVER

Severity	Description
informational	System went up after an upgrade.

IOS_RESTORED_AFTER_HOT_UPGRADE

Severity	Description
informational	System is able to perform I/Os after a hot upgrade.

UPGRADE_NO_NEW_FILES_FOR_UPGRADE

Severity	Description	Troubleshooting
warning	Repository version does not contain any new files. current version *current_version* new version is *new_version*	Contact IBM Support

UPGRADE_DOWNLOAD_REPOSITORY_COPY

Severity	Description	Troubleshooting
critical	Mirroring needed files from repository failed. Mirroring module is *mirroring_module* error is *error*	Contact IBM Support

UPGRADE_LOCAL_VERSION_DOWNLOAD_FAILED

Severity	Description	Troubleshooting
critical	Failure to distribute new sofware internally. Error code is *error* .	Contact IBM Support

UPGRADE_WAS_CANCELLED

Severity	Description	Troubleshooting
informational	Upgrade was cancelled with reason *reason* .	Contact IBM Support

HOT_UPGRADE_ABORTED

Severity	Description	Troubleshooting
critical	Hot upgrade aborted with reason *reason*.	Contact IBM Support

HOT_UPGRADE_HAS_FAILED

Severity	Description	Troubleshooting
critical	Hot upgrade failed while *errorneous_state*.	Contact IBM Support

PRE_UPGRADE_SCRIPT_INVOCATION_FAILED

Severity	Description	Troubleshooting
critical	Invocation of pre-upgrade script failed with error *error*.	Contact IBM Support

POST_UPGRADE_SCRIPT_INVOCATION_FAILED

Severity	Description	Troubleshooting
critical	Invocation of post-upgrade script failed with error *error*.	Contact IBM Support

UPGRADE_IS_NOT_ALLOWED

Severity	Description	Troubleshooting
critical	One or more of the pre-upgrade validations failed.	Fix the problems pointed out it previous events and revalidate.

PRE_UPGRADE_VALIDATION_FAILED

Severity	Description	Troubleshooting
critical	One of the pre-upgrade validations failed with status *error*.	Contact IBM Support

POST_UPGRADE_SCRIPT_STARTED

Severity	Description
informational	Post-upgrade script started.

POST_UPGRADE_SCRIPT_FINISHED

Severity	Description
informational	Post-upgrade script finished successfully.

PRE_UPGRADE_SCRIPT_DISAPPROVES

Severity	Description	Troubleshooting
critical	Upgrade cannot commence because some of the validations in the pre-upgrade script failed. Explanation: *explanation*.	Correct the system state according to the explanation and try again

POST_UPGRADE_SCRIPT_REPORTED_FAILURE

Severity	Description	Troubleshooting
critical	Post upgrade script reported failure. Script output: *explanation*.	Correct the system state according to the explanation and try again

POWER_PROBLEM_CAUSING_MAINTENANCE_MODE

Severity	Description
warning	Power state causing system to enter maintenance mode.

SYSTEM_ENTERED_CHARGING_STATE

Severity	Description
informational	System cannot start work until it is sufficiently charged.

POWER_PROBLEM_CAUSING_MODULE_PHASEOUT

Severity	Description
critical	Module *module* was phased out due to a power problem.

POWER_REPORT_PROBLEM_CAUSING_MODULE_PHASEOUT

Severity	Description
critical	Module *module* did not report power status on time and hence it was phased out.

POWER_PROBLEM_CAUSING_SYSTEM_SHUTDOWN

Severity	Description
critical	Power state causing system to shutdown due to: *Power Emergency Shutdown Reason*.

DELAYING_BACKUP_POWER_FAILURE_HANDLING

Severity	Description

Severity	Description
informational	Delaying backup power failure handling at module *Module*

NO_DELAYED_BACKUP_POWER_FAILURE

Severity	Description
informational	No delayed backup power failure

MODULE_PHASEOUT_FAILED

Severity	Description
informational	Phase out of module *Module* failed

METADATA_SET

Severity	Description
warning	*Object type* with name '*Object name*' has new metadata value.

METADATA_DELETE

Severity	Description
warning	Metadata object deleted for *Object type* with name '*Object name*'.

SUBORDINATE_METADATA_SET

Severity	Description
warning	Remote *Object type* with name '*Object name*' was assigned a new metadata value by local system.

SUBORDINATE_METADATA_DELETE

Severity	Description
warning	Remote metadata object was deleted by local system for *Object type* with name '*Object name*'.

PATCH_SCRIPT_ADDED

Severity	Description	Troubleshooting
informational	Added patch *Patch Name*.	Was patch supposed to have been added.

PATCH_SCRIPT_UPDATED

Severity	Description
informational	Updated patch *Patch Name*.

PATCH_SCRIPT_DELETED

Severity	Description
informational	Deleted patch *Patch Name*.

MODULE_FAILED_TO_FETCH_PATCH_SCRIPT

Severity	Description
warning	Module *Module* failed to fetch patch script *Patch Name*.

PATCH_SCRIPT_FAILED_TO_EXECUTE

Severity	Description
informational	Patch script *Patch Name* execution failed on module *Module*

PATCH_SCRIPT_EXECUTION_STARTED

Severity	Description
informational	Patch script *Patch Name* execution on module *Module* started with pid *Process ID*

PATCH_SCRIPT_EXECUTION_ENDED

Severity	Description
informational	Patch script *Patch Name* execution on module *Module* with pid *Process ID* ended with return code *Return Code*

DOMAIN_CREATED

Severity	Description
informational	Domain *domain_name* has been created.

DOMAIN_UPDATED

Severity	Description
informational	Domain *domain_name* has been updated.

DOMAIN_RENAMED

Severity	Description
informational	Domain *old_name* has been renamed to *domain_name*.

DOMAIN_DELETED

Severity	Description
informational	Domain *domain_name* has been deleted.

POOL_ADDED_TO_DOMAIN

Severity	Description
informational	Pool *pool_name* has been added to domain *domain_name*.

POOL_REMOVED_FROM_DOMAIN

Severity	Description
informational	Pool *pool_name* has been removed from domain *domain_name*.

POOL_MOVED_BETWEEN_DOMAINS

Severity	Description
informational	Pool *pool_name* has been moved from domain *domain_name* to domain *domain_name*.

DOMAINS_AUTO_SHIFT_RESOURCES

Severity	Description
informational	Resources from domain *domain_name* to domain *domain_name* have been auto shifted.

OBJECT_ATTACHED_TO_DOMAIN

Severity	Description
informational	Object *object_name* of type *object_type* has been added to domain *domain_name*.

OBJECT_REMOVED_FROM_DOMAIN

Severity	Description
informational	Object *object_name* of type *object_type* has been removed from domain *domain_name*.

DOMAIN_MANAGED_ATTRIBUTE_SET

Severity	Description
informational	Domain *domain_name* managed attribute was set to *managed_attribute*.

REMOTE_SUPPORT_CONNECTED

Severity	Description
informational	System connected to remote support center *Destination*.

UNABLE_TO_CONNECT_TO_REMOTE_SUPPORT

Severity	Description
minor	System is unable to connect to any remote support center.

REMOTE_SUPPORT_CONNECTION_LOST

Severity	Description
variable	Connection to remote support center *Destination* failed while the connection was in state *Disconnected Session State*.

REMOTE_SUPPORT_TIMEOUT

Severity	Description
variable	Connection to remote support center *Destination* timed out while the connection was in state *Disconnected Session State*.

REMOTE_SUPPORT_IMMINENT_TIMEOUT

Severity	Description
minor	System is about to disconnect busy connection to remote support center *Destination*.

REMOTE_SUPPORT_DEFINED

Severity	Description
informational	Defined remote support center *Name* with IP address *Address* and port *Port*.

REMOTE_SUPPORT_DELETED

Severity	Description
informational	Deleted remote support center *Name*.

REMOTE_SUPPORT_DISCONNECTED

Severity	Description
variable	System disconnected from remote support center *Destination* while the connection was in state *Disconnected Session State*.

REMOTE_SUPPORT_CLIENT_MOVED

Severity	Description
informational	The remote support client moved from *Old Module* to *New Module*.

REMOTE_SUPPORT_CLIENT_NO_AVAILABLE_MODULES

Severity	Description
minor	No live modules with *Port Type* ports are available to run the remote support client.

TIMEZONE_SET

Severity	Description
informational	Timezone of the system was set to *Timezone*.

TIME_SET

Severity	Description	Troubleshooting
informational	On *Previous Time and Date* date and time of the system were set to *Time and Date*.	If date and/or time setting was intended, there is no problem.

TRANSACTION_ROLLED_BACK

Severity	Description
warning	Configuration transaction was rolled back due to module failure. Preceding events may reflect changes that were not committed.

ELICENSE_ACCEPTED

Severity	Description

Severity	Description
informational	Electronic license was accepted by '*Approver Name*'.

ELICENSE_VIOLATION

Severity	Description	Troubleshooting
warning	Latest version of the electronic license was not approved.	Please approve the electronic license.

AUDIT_ENABLED

Severity	Description
informational	CLI command auditing activated.

AUDIT_DISABLED

Severity	Description
warning	CLI command auditing deactivated.

IB_PORT_ENABLE

Severity	Description
informational	Switch port *switch_port* has been enabled.

IB_PORT_DISABLE

Severity	Description
informational	Switch port *switch_port* has been disabled.

IB_PORT_MISWIRE

Severity	Description	Troubleshooting
warning	Switch port miswired: *switch_port* shall connect to *expected_component* but connects to *component* with GUID *guid*.	Contact IBM Support

IB_PORT_MISWIRE_CORRECTED

Severity	Description
informational	Miswire on switch port *switch_port* has been corrected.

IB_PORT_SHUTDOWN

Severity	Description	Troubleshooting
warning	Infiniband port *component* has been shutdown with reason *shutdown_reason*.	Contact IBM Support

IB_PORT_TEST_FAILED

Severity	Description	Troubleshooting
warning	Infiniband port *component* has failed component test with reason *reason*.	Contact IBM Support

IB_PORT_TEST_SUCCESS

Severity	Description	Troubleshooting
informational	Infiniband port *component* completes component test.	Contact IBM Support

IB_PORT_PHASEIN_FAILED

Severity	Description	Troubleshooting
warning	Infiniband port *component* has failed to phase-in with reason *reason*.	Contact IBM Support

IB_PORT_PHASEIN_SUCCESS

Severity	Description	Troubleshooting
informational	Infiniband port *component* completes phase-in.	Contact IBM Support

IB_LINK_DOWN

Severity	Description
warning	Link on *switch_port* (that connects to *component*) is down.

IB_LINK_UP

Severity	Description
informational	Link on *switch_port* (that connects to *component*) is up.

IB_SWITCH_MISSING

Severity	Description	Troubleshooting
major	Switch *ib_switch* is missing.	Contact IBM Support

IB_SWITCH_LOST

Severity	Description	Troubleshooting
warning	Switch *ib_switch* that was missing is considered lost.	Contact IBM Support

IB_MISSING_SWITCH_FOUND

Severity	Description
informational	A previously missing switch *ib_switch* is now found.

IB_CONFIGURE_COMMAND_ERROR

Severity	Description
warning	Infiniband configuration command *command* has failed on *component*.

POD_IB_PORT_MISWIRE

Severity	Description	Troubleshooting
warning	POD module port miswired: *module_port* shall connect to *expected_component* but connects to *component* with GUID *guid*.	Please contact support.

POD_IB_PORT_MISWIRE_CORRECTED

Severity	Description
informational	POD module port *module_port* connected to *component* miswire corrected.

IB_SWITCH_REBOOT_DETECTED

Severity	Description	Troubleshooting
warning	IB switch '*switch_id*' has rebooted.	Contact IBM Support

IB_CONNECTION_SERVICES_UNAVAILABLE

Severity	Description	Troubleshooting
warning	Connection services unavailable on port '*port*'.	Contact IBM Support

IB_CONNECTION_SERVICES_AVAILABLE

Severity	Description	Troubleshooting
informational	Connection services now available on port '*port*'.	Contact IBM Support

REDUCED_IB_PORT_SPEED

Severity	Description	Troubleshooting
major	Component *Component ID* switched to reduced link speed of *Speed* Gbps. Expected link speed is *Expected Speed* Gbps.	Contact IBM Support

IB_PORT_SPEED_OK

Severity	Description	Troubleshooting

Severity	Description	Troubleshooting
informational	Component *Component ID* switched to expected link speed of *Speed* Gbps.	Confirm that link blinking does not take place.

PERF_CLASS_MAX_IO_RATE_UPDATED

Severity	Description
informational	Performance Class *name* max IO rate was changed to *IO rate* IOPS

PERF_CLASS_MAX_BW_RATE_UPDATED

Severity	Description
informational	Performance Class *name* max BW rate was changed to *BW rate* MB/sec

PERF_CLASS_CREATE

Severity	Description
informational	Performance Class with name '*name*' was created

PERF_CLASS_DELETE

Severity	Description
informational	Performance Class with name '*name*' was deleted

PERF_CLASS_ADD_HOST

Severity	Description
informational	Host with name '*host_name*' was added to Performance Class with name '*name*'

PERF_CLASS_REMOVE_HOST

Severity	Description
informational	Host with name '*host_name*' was removed from Performance Class with name '*name*'

PERF_CLASS_ADD_POOL

Severity	Description
informational	Pool with name '*pool.name*' was added to Performance Class with name '*pool.perf_class*'

PERF_CLASS_REMOVE_POOL

Severity	Description
informational	Pool with name '*pool.name*' was removed from Performance Class with name '*name*'

PERF_CLASS_ADD_VOLUME

Severity	Description
informational	Volume with name '*volume.name*' was added to Performance Class with name '*volume.perf_class*'

PERF_CLASS_REMOVE_VOLUME

Severity	Description
informational	Volume with name '*volume.name*' was removed from Performance Class with name '*name*'

PERF_CLASS_ADD_DOMAIN

Severity	Description
informational	Domain *domain_name* was added to Performance Class *name*

PERF_CLASS_REMOVE_DOMAIN

Severity	Description
informational	Domain *domain_name* was removed from Performance Class *name*

VOLUME_MODIFIED_DURING_IO_PAUSE

Severity	Description	Troubleshooting
warning	Volume '*vol_name*' of CG '*cg_name*' was modified during Pause IO with token '*token*'	Retry after completing CG changes.

CONS_GROUP_MODIFIED_DURING_IO_PAUSE

Severity	Description	Troubleshooting
warning	CG '*cg_name*' was modified during Pause IO with token '*token*'	Retry after completing CG changes.

IO_PAUSED_FOR_CONS_GROUP

Severity	Description
Informational	Pause IO on CG with name '*cg_name*' was started with *timeout*ms timeout . Token is '*token*'.

IO_RESUMED_FOR_CONS_GROUP_EXPLICITLY

Severity	Description
Informational	Pause IO on CG with name '*cg_name*' and token '*token*' was resumed by user request.

IO_RESUMED_FOR_CONS_GROUP_AUTOMATICALLY

Severity	Description
Informational	Pause IO on CG with name '*cg_name*' and token '*token*' was resumed after snapgroup creation.

IO_RESUMED_FOR_CONS_GROUP_UPON_SYSTEM_ERROR

Severity	Description
warning	Pause IO on CG with name '*cg_name*' and token '*token*' was resumed after system error.

IO_RESUMED_FOR_CONS_GROUP_UPON_TIMEOUT_EXPIRATION

Severity	Description	Troubleshooting
warning	Pause IO on CG with name '*cg_name*' and token '*token*' was canceled after timeout.	Use longer timeout value or require less time for performing action.

ALU_CREATE

Severity	Description
informational	ALU was defined with name '*ALU name*' associated with host '*ALU host name*' lun '*ALU lun*'.

ALU_DELETE

Severity	Description
informational	ALU with name '*ALU name*' associated with host '*ALU host name*' lun '*ALU lun*' was deleted.

ALU_UNBOUND_ALL

Severity	Description
informational	All SLUs of ALU with name '*ALU name*' were unbound.

TXN_NODE_FLASH_CONNECTION_LOST

Severity	Description
warning	TXN node on '*module*' is disconnected from FlashSystem '*FlashSystem*'.

TXN_NODE_FLASH_CONNECTED

Severity	Description
informational	TXN node on '*module*' is fully connected to FlashSystem '*FlashSystem*'.

POD_IB_LINK_DETECTION_LINK_PERSISTENTLY_DISCONNECTED

Severity	Description

Severity	Description
critical	IB link from *Source* to *Target* has reported as persistently disconnected

POD_IB_LINK_DETECTION_LINK_PERSISTENTLY_CONNECTED

Severity	Description
informational	IB link from *Source* to *Target* has detected as persistently connected

Return Codes

The following table specifies the return codes.

Return Code	Error Description
0	Success.
1	Command execution failed.
2	No connection to the system.
3	Password is required.
4	Password does not match system password.
7	Command not allowed from this client.
8	Bad XCLI option.
9	Internal XCLI error.

REST API reference

The IBM Hyper-Scale Manager REST Application Programming Interface (API) is documented within the IBM Hyper-Scale Manager documentation set.

For REST API product information and documentation, see IBM Hyper-Scale Manager documentation.

Open API reference

The following information describes how to use the Open Application Programming Interface (API) for IBM FlashSystem® A9000 and A9000R systems.

Introduction

The IBM® FlashSystem A9000 and A9000R Open Application Programming Interface (API) is a non-proprietary storage-management client application.

The Open API uses the Storage Management Initiative Specification (SMI-S), as defined by the Storage Networking Industry Association (SNIA) to view LUN information.

You can use the Open API to integrate configuration-management support into storage resource management (SRM) applications so that you can use your existing SRM application and infrastructures to configure and manage IBM FlashSystem® A9000 and A9000R systems. The Open API presents another option for managing your systems by complementing the use of IBM Hyper-Scale Manager and the command-line interface (CLI). The Open API is an embedded component of IBM FlashSystem A9000 and A9000R systems.

You can implement the Open API without using a separate middleware application, like the IBM System Storage Common Information Model (CIM) agent, which provides a CIM-compliant interface. The Open API uses the CIM technology to manage proprietary devices as open system devices through storage management applications. The Open API is used by storage management applications to communicate with IBM FlashSystem A9000 and A9000R systems.

- **Storage Management Initiative Specifications (SMI-S) overview**
  The Storage Management Initiative Specifications (SMI-S) are the leading storage-management API standards that are widely adopted by storage vendors. Using the SMI-S API, customers or independent software vendors (ISVs) can develop their own software solutions for managing heterogeneous storage deployments. For example, Microsoft is using SMI-S in Microsoft System Center Virtual Machine Manager 2012 to provide rapidly-integrated management of advanced disk arrays without needing any storage vendor add-on.
- **Common Information Model (CIM) agent overview**
  A Common Information Model (CIM) agent provides a means by which a device can be managed by common building blocks rather than proprietary software. If a device is CIM-compliant, software that is also CIM-compliant can manage the device. Vendor applications can manage CIM-compliant devices in a common way, rather than using device-specific programming interfaces. You can perform tasks in a consistent manner across devices and vendor applications.
- **Common Information Model (CIM) agent components**
  With a Common Information Model (CIM) agent, application programmers can use common building blocks rather than proprietary software or device-specific programming interfaces to manage CIM-compliant devices. Standardization of the way that applications manage storage provides easier storage management.
- **Open API element definitions**
  The IBM FlashSystem A9000 and A9000R Open API elements include schemas, classes, properties, methods, indications, associations, references, and

qualifiers, as described in the following list.

- **Common Information Model (CIM) agent security**
  The Common Information Model (CIM) agent can operate in both secure and unsecure modes.

Storage Management Initiative Specifications (SMI-S) overview

The Storage Management Initiative Specifications (SMI-S) are the leading storage-management API standards that are widely adopted by storage vendors. Using the SMI-S API, customers or independent software vendors (ISVs) can develop their own software solutions for managing heterogeneous storage deployments. For example, Microsoft is using SMI-S in Microsoft System Center Virtual Machine Manager 2012 to provide rapidly-integrated management of advanced disk arrays without needing any storage vendor add-on.

All IBM storage products support the SMI-S standard, which defines the Common Information Model (CIM) of each storage system management and is used as the base model for ISV and customer solution development. The built-in support for SMI-S in IBM storage systems, including IBM FlashSystem A9000 and A9000R, facilitates the development of different management solutions such as monitoring and control of physical storage, configuration and provisioning of SAN storage, and copying or protection of data through remote mirroring and snapshots.

The SMI-S standard is driven by the Storage Networking Industry Association (SNIA), and the CIM standard is driven by the Distributed Management Task Force (DMTF).

Common Information Model (CIM) agent overview

A Common Information Model (CIM) agent provides a means by which a device can be managed by common building blocks rather than proprietary software. If a device is CIM-compliant, software that is also CIM-compliant can manage the device. Vendor applications can manage CIM-compliant devices in a common way, rather than using device-specific programming interfaces. You can perform tasks in a consistent manner across devices and vendor applications.

A CIM agent consists of the components that are shown in Figure 1. The main components are the CIM object manager (CIMOM), the service location protocol (SLP), and the device provider. A device can be a storage system such as IBM FlashSystem A9000 or IBM FlashSystem A9000R. The CIM agent registers itself with the SLP Service Agent (SLP SA) to enable discovery by the client application.

The SLP DA is a directory service daemon that a client application calls to locate the CIMOM. The client application and the CIMOM communicate through CIM messages. The CIMOM and device provider communicate through method calls made from the CIMOM to the provider. The device provider communicates with the device through proprietary calls.

Figure 1. How a CIM agent works

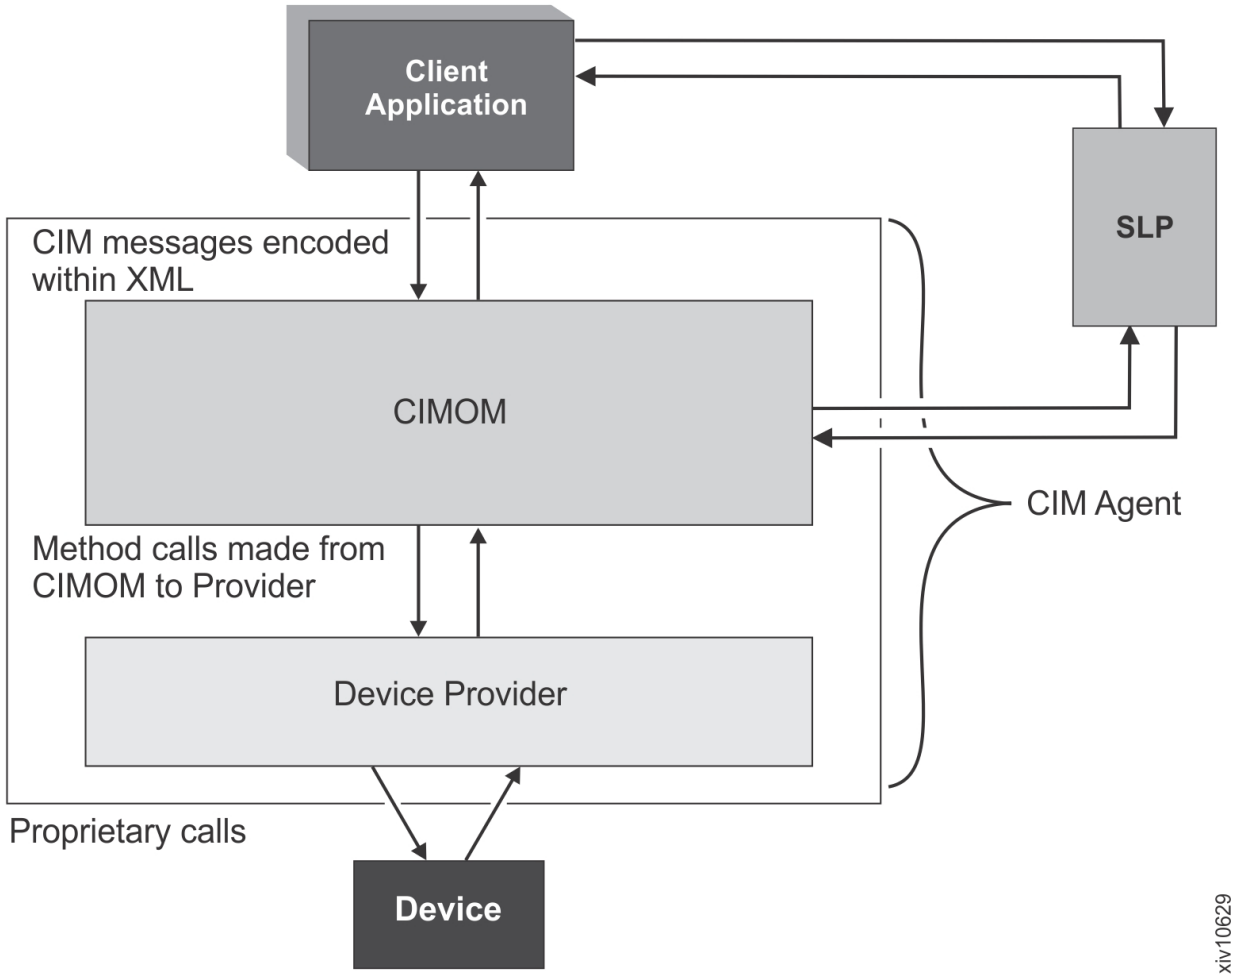

The CIMOM supports the following specifications and standards:

- Distributed Management Task Force (DMTF) Specification for CIM Operations over HTTP, Version 1.4
- CIM Specification, version 2.40.0
- Storage Networking Industry Association (SNIA) Storage Management Initiative Specification (SMI-S) and the Shared Storage Model (SSM), a framework for describing storage architectures, version 1.6

These specifications allow a CIM agent to act as an open-system standards interpreter. Other CIM-compliant storage resource management applications (IBM® and non-IBM) can interoperate with each other.
When you install, configure, and enable the CIM agent on a host server or an administrative workstation within your network, that host server or workstation can communicate with your IBM FlashSystem A9000 or IBM FlashSystem A9000R through the CIM agent. CIM-compliant applications like the CIM agent can manage the data on your system.

The following sites provide more information about the CIM standards:

- DMTF Common Information Model (CIM) Standards
- Storage Networking Industry Association Standards

- **Common Information Model (CIM) concepts**
  The Common Information Model (CIM) is an open approach to the management of systems and networks.

Common Information Model (CIM) concepts

The Common Information Model (CIM) is an open approach to the management of systems and networks.

The CIM provides a common conceptual framework applicable to all areas of management, which includes systems, applications, databases, networks, and devices. The CIM specification provides the language and the methodology that is used to describe management data.

The CIM defines a set of classes with properties and associations that provide a conceptual framework. The framework can be used to organize data for a specific managed environment, such as data storage. CIM Schema 2.11 for Managing a Storage Array provides information about enabling management applications to manage data in a common way.

The CIM standards and the Distributed Management Task Force (DMTF) specification provide information about Web-based enterprise management (WBEM) operations over HTTP.

When the CIM object manager (CIMOM) first starts, it registers itself to the SLP. It provides information about its location (IP address and port) and the type of service it provides. A client application finds the location of the CIMOM by calling an SLP directory service. After this information is obtained, the client application opens direct communication with the CIMOM.

A client sends requests to a CIMOM in the context of a CIM model. The model is defined by the CIM schema and loaded into the repository of the CIMOM. Figure 1 shows how the schema is loaded into the data store of the CIMOM. The managed object format (MOF) compilation and creation of the data store is managed automatically during installation.

As requests arrive, the CIMOM validates and authenticates each request. Requests are either directed to the appropriate functional component of the CIMOM or directed to a device-specific handler called a provider.

Figure 1. The MOF compiler stores the model in the CIMOM data store.

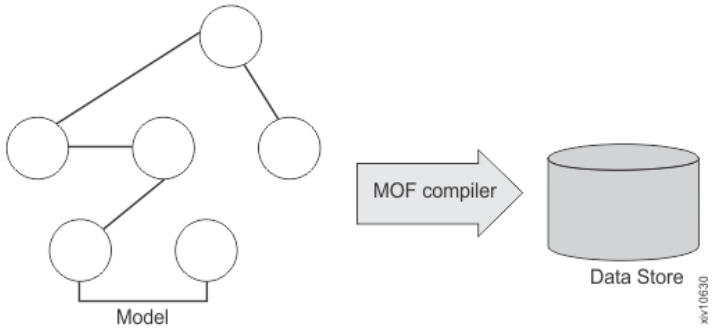

A provider makes device-unique programming interface calls on behalf of the CIMOM to satisfy a client application request. Such requests generally map a CIM request to the API for a device. A request to get an instance of a class or a property of an instance, for example, might be directed to a provider. Then, the provider might make one or many requests of a device by using the unique API for the device. Figure 1 shows the communication structure between the device and the client application.

Common Information Model (CIM) agent components

With a Common Information Model (CIM) agent, application programmers can use common building blocks rather than proprietary software or device-specific programming interfaces to manage CIM-compliant devices. Standardization of the way that applications manage storage provides easier storage management.

The following list describes a CIM agent and its components:

CIM agent
: An agent that interprets open-system data as it is transferred between the API and a device or a storage unit.

CIM object manager (CIMOM)
: A common conceptual framework for data management. Receives, validates, and authenticates client application requests, and then directs requests to the appropriate functional component or to a device provider.

client application
: A storage management API that initiates a request to a device or a data storage unit such as an IBM FlashSystem A9000 or IBM FlashSystem A9000R.
Note: A client application is not provided with the CIM agent, and it must be supplied by the customer

Service Location Protocol (SLP)
: The SLP DA is a directory service that a client application calls to locate the CIMOM. The SLP SA is a service agent that enables discovery by a client application.

storage unit (also known as a storage server)
: The final destination of a client application request and the processor of the request.

storage unit provider
: A storage unit-specific handler that receives client application requests that are destined for its device or storage unit.

Open API element definitions

The IBM FlashSystem A9000 and A9000R Open API elements include schemas, classes, properties, methods, indications, associations, references, and qualifiers, as described in the following list.

Schema
: A group of classes that are defined to a single namespace. Within the Common Information Model (CIM) agent, the schemas that are supported are the ones that are loaded through the managed object format (MOF) compiler.

Class
: The definition of an object within some hierarchy. Classes can have methods and properties and be the target of an association.

Property
: A value that is used to characterize instances of a class.

Method
: An implementation of a function on a class.

Indication
: An object representation of an event.

Association
: A class that contains two references which define a relationship between two objects.

Reference
: A unique identifier of an object that is based on its key properties.

Qualifier
: Additional information about other elements, classes, associations, indications, methods, method parameters, instances, properties, or references.

Namespace

A namespace defines the scope over which an IBM FlashSystem A9000 and A9000R Open API schema applies. IBM FlashSystem A9000 and A9000R Open API operations are always run within the context of a namespace. An IBM FlashSystem A9000 and A9000R Open API schema or version is loaded into a namespace when that schema is compiled by the MOF compiler. The namespace must be specified within the message that the client sends to the Open API.

Clients cannot create new namespaces. Attempts to do so result in errors.

Object name

An object name consists of a namespace path and a model path. The namespace path provides access to the API implementation that is managed by the IBM FlashSystem A9000 and A9000R Open API. The model path provides navigation within the implementation. The following example shows an object name:

```
http://cimom.host.com/root/ibm:CIM_Class.key1=value1,key2=value2
```

where `http://cimom.host.com/root/ibm` is the namespace path and `:CIM_Class.key1=value1,key2=value2` is the model path.

Common Information Model (CIM) agent security

The Common Information Model (CIM) agent can operate in both secure and unsecure modes.

Secure mode
: All requests between the client application and the CIM object manager (CIMOM) are XML encoded requests. They are sent over HTTP or HTTP over Secure Sockets Layer (SSL). The CIMOM, upon receiving a request, parses the request and processes it. Responses, when they are returned to the client application, are transformed into XML-encoded CIM status and returned in HTTP responses to the client. The CIM agent runs in secure mode by using SSL by default.

Unsecure mode
: Some vendor software cannot communicate with the CIM agent in a secure mode. You can still use this vendor software by configuring the CIM agent to run with only basic user name and password security. See the configuration instructions for your operating system for the instructions for configuring the CIM agent for this less secure mode.

Installation and configuration

The Common Information Model (CIM) agent is preinstalled and embedded on the administrative module, and it is enabled by default and preconfigured. You can manage IBM FlashSystem® A9000 and A9000R systems from the CIM agent that is bundled with the administrative module.

Additional details regarding the preinstalled and preconfigured CIM agent:

- The embedded CIM agent agent does not require configuration changes to manage IBM FlashSystem A9000 and A9000R systems.
- The CIM object manager (CIMOM) is installed and running on all three administrative modules. The clients can connect to any of the administrative modules, and the same results are provided to the CIMOM.
- You can use the cim_enable and cim_disable commands through the command-line interface (CLI) to enable or disable the CIM agent. While the CIM agent is enabled, a watchdog process monitors it to ensure it is always running.

The CIM agent embedded on the administrative module has the following limitations:

- The CIM agent can support only IBM FlashSystem A9000 and A9000R systems on which the administrative module is located. The CIM agent is not able to manage any other IBM FlashSystem A9000 or IBM FlashSystem A9000R system.
- The CIM agent must use secure connections over port 5989.
- The CIM agent uses the IBM FlashSystem A9000 or IBM FlashSystem A9000R system user account to authenticate. To manage accounts, you must use IBM® Hyper-Scale Manager or the CLI.

Note: CIM uses XML format when communicating with the IBM FlashSystem A9000 and A9000R systems, so the user name and password can't include the following special characters: `&`, `<`, `>`, `'`, `"`
If you must use these characters in your user name and password, manually replace them with their XML equivalents, as shown below:

```
& => &amp;

< => &lt;

> => &gt;

' => &apos;

" => &quot;
```

- **Multitenancy feature support information**
  For IBM FlashSystem A9000 and A9000R, the multitenancy feature is supported.

Multitenancy feature support information

For IBM FlashSystem® A9000 and A9000R, the multitenancy feature is supported.

To support this feature, IBM FlashSystem A9000 and A9000R implement the concept of a domain that will link users to their dedicated pools, hosts, mirror targets, and other resources. A user can manage only the storage resources to which they are associated, without the ability to modify or monitor other system resources. The domain restricts the set of objects a user can manage to those associated with the domain. If a user is associated with one or more domains, they are a domain user, otherwise, they are a global user. This means that users in the same category but different domains will manage different resources.

The Common Information Model (CIM) agent for IBM FlashSystem A9000 and A9000R, however, does *not* support the multitenacy feature. CIM only supports the global user with the domain policy access value of OPEN. CIM requires a global Storage administrator/Application administrator/Read-only user who can manage all the resources in the system.

See IBM FlashSystem A9000 and A9000R: Architecture, Implementation, and Usage and the *IBM FlashSystem A9000 Product Overview* or *IBM FlashSystem A9000R Product Overview* for more information about the architecture of IBM FlashSystem A9000 and A9000R systems.

Communication concepts and methods

Refer to the following topics for information about Common Information Model (CIM) agent communication concepts and methods.

- **Common Information Model (CIM) agent communication concepts**
  The Common Information Model (CIM) agent uses client communication and intrinsic and extrinsic methods of communication.
- **Common Information Model (CIM) agent communication methods**
  Client application calls to intrinsic methods can result in Common Information Model (CIM) agent calls to the device provider. This result happens when the device provider surfaces the classes or instances that are referenced in the calls.
- **Return error codes**
  The Common Information Model object manager (CIMOM) returns status to the client application.

Common Information Model (CIM) agent communication concepts

The Common Information Model (CIM) agent uses client communication and intrinsic and extrinsic methods of communication.

Client communication

A client application communicates with the CIM agent through operation request messages that are encoded within XML. The CIM agent returns responses with operation response messages. Requests and responses are sub-elements of the **<CIM MESSAGE>** element.

A **<MESSAGE>** sent to the CIM agent must contain an ID attribute. A response from the CIM agent returns this value and enables the client to track requests and their responses.

The CIM agent supports simple requests and simple responses. *Simple requests* are operation request messages that contain the **<SIMPLEREQ>** XML tag. *Simple responses* are operation response messages that contain the **<SIMPLERSP>** XML tag. A client application determines that the CIM agent supports only simple operation requests and responses by examining the output of the OPTIONS method.

Intrinsic and extrinsic methods

All operations on the CIM agent are completed by running one or more methods. A method is either an intrinsic method or an extrinsic method.

Intrinsic methods are supported by the CIM agent itself. These methods are included within **<IMETHODCALL>** XML tags that are sent in messages to the CIM agent.

Extrinsic methods are defined by the schema that is supported by the CIM agent. These methods are included within **<METHODCALL>** XML tags that are sent in messages to the CIM agent. Client applications can call on the CIM agent by using these methods. These methods fall within certain functional groups that might be supported by the CIM agent.

Common Information Model (CIM) agent communication methods

Client application calls to intrinsic methods can result in Common Information Model (CIM) agent calls to the device provider. This result happens when the device provider surfaces the classes or instances that are referenced in the calls.

The CIM agent returns **IMETHODRESPONSE** or **METHODRESPONSE** elements to the client application when the intrinsic or extrinsic methods are used. These elements are contained within a **MESSAGERESPONSE** XML tag.

Functional groups

Table 1 describes the functional groups that are supported by the CIM agent. This information is also returned to the client that makes an OPTIONS request to the CIM agent.

Table 1. Functional groups for the CIM agent

Functional group	Parameters	Supported or not supported
Basic read	GetInstance EnumerateInstances EnumerateInstanceNames GetProperty	Supported
Basic write	SetProperty	Not Supported
Schema manipulation	CreateClass ModifyClass DeleteClass	Not Supported
Instance manipulation	CreateInstance ModifyInstance DeleteInstance	Not Supported
Association traversal	Associators AssociatorNames References RefernceNames	Supported
Qualifier read	GetQualifier EnumerateQualifiers	Supported
Qualifier manipulation	SetQualifier DeleteQualifier	Not Supported
Query execution	ExecQuery	Supported

The most current information for the communication methods is in the MOF documentation. The MOF documentation is in the mof folder in the CIM agent installation directory.

- **Obtaining a single class from the target namespace**
  The GetClass method returns a single class from the target namespace.
- **Obtaining a single instance from the target namespace**
  The GetInstance method returns a single instance from the target namespace.
- **Deleting a single instance from the target namespace**
  The DeleteInstance method deletes a single instance from the target namespace.
- **Creating an instance in the target namespace**
  The CreateInstance method creates an instance in the target namespace. To use this method, the instance cannot exist.
- **Modifying an existing instance in the target namespace**
  The ModifyInstance method modifies an existing instance in the target namespace. The instance must exist.
- **Enumerating classes within a defined target namespace**
  The EnumerateClasses method enumerates classes within a defined target namespace.
- **Enumerating the names of subclasses of a class defined within the target namespace**
  The EnumerateClassNames method enumerates the names of subclasses of a class that is defined within the target namespace.
- **Enumerating instances of a defined class within the target namespace**
  The EnumerateInstances method enumerates instances of a defined class within the target namespace.
- **Enumerating the names of instances of a class within a target namespace**
  The EnumerateInstanceNames method enumerates the names of the instances of a class within a target namespace.
- **Processing a query against the target namespace**
  The ExecuteQuery method processes a query against the target namespace.
- **Enumerating classes or instances that are associated with a specific Common Information Model (CIM) object**
  The Associators method enumerates classes or instances that are associated with a specific Common Information Model (CIM)object.
- **Enumerating the names of classes or instances associated with a specific Common Information Model (CIM) object**
  The AssociatorNames method enumerates the names of the classes or instances that are associated with a specific Common Information Model (CIM) object.
- **Enumerating the association objects that refer to a specific target class or instance**
  The References method enumerates the association objects that refer to a particular target class or instance.
- **Enumerating the names of the association objects that refer to a specific target class or instance**
  The ReferenceNames method enumerates the names of the association objects that refer to a particular target class or instance.
- **Retrieving a single property value from an instance in the target namespace**
  The GetProperty method retrieves a single property value from an instance in the target namespace.
- **Setting a single property value within an instance in the target namespace**
  The SetProperty method sets a single property value within an instance in the target namespace. But the IBM FlashSystem A9000 and A9000R Open API Common Information Model (CIM) agent does not have any features that use this method.
- **Retrieving a single qualifier declaration from the target namespace**
  The GetQualifier method retrieves a single qualifier declaration from the target namespace.
- **Creating or modifying a qualifier declaration in the target namespace**
  The SetQualifier method creates or updates a qualifier declaration in the target namespace.
- **Enumerating qualifier declarations from the target namespace**
  The EnumerateQualifiers method enumerates qualifier declarations from the target namespace.
- **Common Information Model (CIM) agent communication methods that cannot be used**
  The following Common Information Model (CIM) agent communication methods cannot be used because they are not supported by IBM FlashSystem® A9000 and A9000R:

Obtaining a single class from the target namespace

The GetClass method returns a single class from the target namespace.

Parameters

Table 1 describes the parameters of the GetClass method.

Table 1. GetClass method parameters

Parameter	Type	Description
ClassName	String	Defines the name of the class you want to retrieve.
LocalOnly	Boolean	TRUE returns all properties, methods, and qualifiers that are overridden within the definition of the class.
IncludeQualifiers	Boolean	TRUE returns all qualifiers for the class, its properties, methods, or method parameters. FALSE returns no qualifiers.
IncludeClassOrigin	Boolean	TRUE returns the CLASSORIGIN attribute of the class.

Return values

If successful, a single class is returned. Otherwise, one of the following error codes is returned:

- CIM_ERR_ACCESS_DENIED
- CIM_ERR_INVALID_NAMESPACE
- CIM_ERR_INVALID_PARAMETER
- CIM_ERR_FAILED

Obtaining a single instance from the target namespace

The GetInstance method returns a single instance from the target namespace.

Parameters

Table 1 describes the parameters of the GetInstance method.

Table 1. GetInstance method parameters

Parameter	Type	Description
InstanceName	String	Defines the name of the instance to retrieve.
IncludeClassOrigin	Boolean	TRUE returns the CLASSORIGIN attribute of the class.

Return values

If successful, a single class is returned. Otherwise, one of the following error codes is returned:

- CIM_ERR_ACCESS_DENIED
- CIM_ERR_INVALID_NAMESPACE
- CIM_ERR_INVALID_PARAMETER
- CIM_ERR_INVALID_CLASS
- CIM_ERR_NOT_FOUND
- CIM_ERR_FAILED

Deleting a single instance from the target namespace

The DeleteInstance method deletes a single instance from the target namespace.

Parameters

Table 1 describes the parameters of the DeleteInstance method.

Table 1. DeleteInstance method parameters

Parameter	Type	Description
InstanceName	String	Defines the name of the instance you want to delete.

Return values

The named instance is deleted or one of the following error codes is returned:

- CIM_ERR_ACCESS_DENIED

- CIM_ERR_INVALID_NAMESPACE
- CIM_ERR_INVALID_PARAMETER
- CIM_ERR_INVALID_CLASS
- CIM_ERR_NOT_FOUND
- CIM_ERR_FAILED

Creating an instance in the target namespace

The CreateInstance method creates an instance in the target namespace. To use this method, the instance cannot exist.

The CreateInstance method is a standard Common Information Model (CIM) method. The IBM FlashSystem A9000 and A9000R Open API does not have any features that use this method.

Parameters

Table 1 describes the parameters of the CreateInstance method.

Table 1. CreateInstance method parameters

Parameter	Type	Description
Instance	Object	The instance to be created. The instance must be based on a class that is already defined in the target namespace.

Return values

If successful, the specified instance is created. Otherwise, one of the following error codes is returned:

- CIM_ERR_ACCESS_DENIED
- CIM_ERR_INVALID_NAMESPACE
- CIM_ERR_INVALID_PARAMETER
- CIM_ERR_INVALID_CLASS
- CIM_ERR_ALREADY_EXISTS
- CIM_ERR_FAILED

Modifying an existing instance in the target namespace

The ModifyInstance method modifies an existing instance in the target namespace. The instance must exist.

Parameters

Table 1 describes the parameters of the ModifyInstance method.

Table 1. ModifyInstance method parameters

Parameter	Type	Description
Instance	Object	Defines the modified instance.

Return values

If successful, the specified instance is updated. Otherwise, one of the following error codes is returned:

- CIM_ERR_ACCESS_DENIED
- CIM_ERR_INVALID_NAMESPACE
- CIM_ERR_INVALID_PARAMETER
- CIM_ERR_INVALID_CLASS
- CIM_ERR_NOT_FOUND
- CIM_ERR_FAILED

Enumerating classes within a defined target namespace

The EnumerateClasses method enumerates classes within a defined target namespace.

Parameters

Table 1 describes the parameters of the EnumerateClasses method.

Table 1. EnumerateClasses method parameters

Parameter	Type	Description

Parameter	Type	Description
ClassName	String	Defines the name of the class for which subclasses are to be returned. If this field is NULL, all base classes within the target namespace are returned.
DeepInheritance	Boolean	TRUE returns all subclasses of the specified class. FALSE returns only immediate child subclasses.
LocalOnly	Boolean	TRUE returns all properties, methods, and qualifiers, that are overridden within the definition of the class.
IncludeQualifiers	Boolean	TRUE returns all qualifiers for the class, its properties, methods, or method parameters. FALSE returns no qualifiers.
IncludeClassOrigin	Boolean	TRUE returns the CLASSORIGIN of the class.

Return values

If successful, zero or more classes (CIMClass) are returned. Otherwise, one of the following error codes is returned:

- CIM_ERR_ACCESS_DENIED
- CIM_ERR_INVALID_NAMESPACE
- CIM_ERR_INVALID_PARAMETER
- CIM_ERR_INVALID_CLASS
- CIM_ERR_FAILED

Enumerating the names of subclasses of a class defined within the target namespace

The EnumerateClassNames method enumerates the names of subclasses of a class that is defined within the target namespace.

Parameters

Table 1 describes the parameters of the EnumerateClassNames method.

Table 1. EnumerateClassNames method parameters

Parameter	Type	Description
ClassName	String	Defines the name of the class for which subclass names are to be returned. If this field is NULL, all base class names within the target namespace are returned.
DeepInheritance	Boolean	TRUE returns all subclass names of the specified class. FALSE returns only immediate child subclass names.

Return values

If successful, zero or more class names are returned. Otherwise, one of the following error codes is returned:

- CIM_ERR_ACCESS_DENIED
- CIM_ERR_INVALID_NAMESPACE
- CIM_ERR_INVALID_PARAMETER
- CIM_ERR_INVALID_CLASS
- CIM_ERR_FAILED

Enumerating instances of a defined class within the target namespace

The EnumerateInstances method enumerates instances of a defined class within the target namespace.

Parameters

Table 1 describes the parameters of the EnumerateInstances method.

Table 1. EnumerateInstances method parameters

Parameter	Type	Description
ClassName	String	Defines the name of the class for which instances are to be returned.
DeepInheritance	Boolean	TRUE returns all instances and all properties of the instance, including the properties added by subclassing. FALSE returns only properties that are defined for the specified class.
IncludeClassOrigin	Boolean	TRUE returns the CLASSORIGIN attribute of the class within the instance.

Return values

If successful, zero or more instances (Objects) are returned. Otherwise, one of the following error codes is returned:

- CIM_ERR_ACCESS_DENIED
- CIM_ERR_INVALID_NAMESPACE
- CIM_ERR_INVALID_PARAMETER
- CIM_ERR_INVALID_CLASS
- CIM_ERR_FAILED

Enumerating the names of instances of a class within a target namespace

The EnumerateInstanceNames method enumerates the names of the instances of a class within a target namespace.

Parameters

Table 1 describes the parameter of the EnumerateInstanceNames method.

Table 1. EnumerateInstanceNames method parameters

Header	Header	Description
ClassName	String	Defines the name of the class for which instance names are returned.

Return values

If successful, zero or more names of instances are returned. Otherwise, one of the following error codes is returned:

- CIM_ERR_ACCESS_DENIED
- CIM_ERR_INVALID_NAMESPACE
- CIM_ERR_INVALID_PARAMETER
- CIM_ERR_INVALID_CLASS
- CIM_ERR_FAILED

Processing a query against the target namespace

The ExecuteQuery method processes a query against the target namespace.

Parameters

Table 1 describes the parameters of the ExecuteQuery method.

Table 1. ExecuteQuery method parameters

Parameter	Type	Description
QueryLanguage	String	Defines the query language in which the query parameter is expressed.
Query	String	Defines the query to be initiated.

Return values

If successful, the method returns a table definition, followed by zero or more rows that correspond to the results of the query. Otherwise, one of the following error codes is returned:

- CIM_ERR_ACCESS_DENIED
- CIM_ERR_NOT_SUPPORTED
- CIM_ERR_INVALID_NAMESPACE
- CIM_ERR_INVALID_PARAMETER
- CIM_ERR_QUERY_LANGUAGE_NOT_SUPPORTED
- CIM_ERR_QUERY_FEATURE_NOT_SUPPORTED
- CIM_ERR_INVALID_QUERY
- CIM_ERR_FAILED

Enumerating classes or instances that are associated with a specific Common Information Model (CIM) object

The Associators method enumerates classes or instances that are associated with a specific Common Information Model (CIM)object.

Parameters

Table 1 describes the parameters of the Associators method.

Table 1. Associators method parameters

Parameter	Type	Description
ObjectName	String	Defines the class name or instance name that is the source of the association.
AssocClass	String	If not NULL, indicates that all objects must be associated with the source object through an instance of this class or one of its subclasses.
ResultClass	String	If not NULL, indicates that all returned objects must be instances of this class or one of its subclasses or be this class.
Role	String	If not NULL, indicates that each return object must be associated with the source object that plays the specified role. The name of the property in the association class that refers to the source object must match the value of this parameter.
ResultRole	String	If not NULL, indicates that each returned object must be associated with the source object that plays the specified role. That is, the name of the property in the association class that refers to the returned object must match the value of this parameter.
IncludeClassOrigin	Boolean	TRUE returns the CLASSORIGIN attribute of the class.

Return values

If successful, zero or more classes (CIMClass) or instances (Objects) are returned. Otherwise, one of the following error codes is returned:

- CIM_ERR_ACCESS_DENIED
- CIM_ERR_INVALID_NAMESPACE
- CIM_ERR_INVALID_PARAMETER
- CIM_ERR_INVALID_CLASS
- CIM_ERR_FAILED

Enumerating the names of classes or instances associated with a specific Common Information Model (CIM) object

The AssociatorNames method enumerates the names of the classes or instances that are associated with a specific Common Information Model (CIM) object.

Parameters

Table 1 describes the parameters of the AssociatorNames method.

Table 1. Associators method parameters

Parameter	Type	Description
ObjectName	String	Defines the class name or instance name that is the source of the association.
AssocClass	String	If not NULL, indicates that returned objects are associated with the source object through an instance of this class or one of its subclasses.
ResultClass	String	If not NULL, indicates that all returned object paths must identify instances of this class or one of its subclasses or must be this class.
Role	String	If not NULL, the name of the property in the association class that refers to the source object must match the value of this parameter.
ResultRole	String	If not NULL, the name of the property in the association class that refers to the return object must match the value of this parameter.

Return values

If successful, zero or more class paths (CIMObjectPath) are returned. Otherwise, one of the following error codes is returned:

- CIM_ERR_ACCESS_DENIED
- CIM_ERR_INVALID_NAMESPACE
- CIM_ERR_INVALID_PARAMETER
- CIM_ERR_FAILED

Enumerating the association objects that refer to a specific target class or instance

The References method enumerates the association objects that refer to a particular target class or instance.

Parameters

Table 1 describes the parameters of the References method.

Table 1. References method parameters

Parameter	Type	Description
ObjectName	String	Defines the class name or instance name whose referring objects are to be returned.
ResultClass	String	If not NULL, indicates that all returned objects must be instances of this class or one of its subclasses or must be this class.
Role	String	If not NULL, must be a valid property name. Each returned object must refer to the target object through a property whose name matches the value of this parameter.
IncludeClassOrigin	Boolean	TRUE returns the CLASSORIGIN attribute of the class.

Return values

If successful, zero or more classes (CIMClass) or instances (Objects) are returned. Otherwise, one of the following error codes is returned:

- CIM_ERR_ACCESS_DENIED
- CIM_ERR_INVALID_NAMESPACE
- CIM_ERR_INVALID_PARAMETER
- CIM_ERR_INVALID_CLASS
- CIM_ERR_FAILED

Enumerating the names of the association objects that refer to a specific target class or instance

The ReferenceNames method enumerates the names of the association objects that refer to a particular target class or instance.

Parameters

Table 1 describes the parameters of the ReferenceNames method.

Table 1. ReferenceNames method parameters

Parameter	Type	Description
ObjectName	String	Defines the class name or instance name whose referring objects are to be returned.
ResultClass	String	If not NULL, indicates that all returned objects must be instances of this class or one of its subclasses or must be this class.
Role	String	If not NULL, must be a valid property name. Each returned object must refer to the target object through a property whose name matches the value of this parameter.

Return values

If successful, the return value specifies the value of the requested property. Otherwise, one of the following error codes is returned:

- CIM_ERR_ACCESS_DENIED
- CIM_ERR_INVALID_NAMESPACE
- CIM_ERR_INVALID_PARAMETER
- CIM_ERR_INVALID_CLASS
- CIM_ERR_NOT_FOUND
- CIM_ERR_NO_SUCH_PROPERTY
- CIM_ERR_FAILED

Retrieving a single property value from an instance in the target namespace

The GetProperty method retrieves a single property value from an instance in the target namespace.

Parameters

Table 1 describes the parameters of the GetProperty method.

Table 1. GetProperty method parameters

Parameter	Type	Description
InstanceName	String	Defines the name of the instance.
Property	String	The name of the property whose value is to be returned from the instance.

Return values

If successful, the return value specifies the value of the requested property. Otherwise, one of the following return codes is returned:

- CIM_ERR_ACCESS_DENIED
- CIM_ERR_INVALID_NAMESPACE

- CIM_ERR_INVALID_PARAMETER
- CIM_ERR_INVALID_CLASS
- CIM_ERR_NOT_FOUND
- CIM_ERR_NO_SUCH_PROPERTY
- CIM_ERR_FAILED

Setting a single property value within an instance in the target namespace

The SetProperty method sets a single property value within an instance in the target namespace. But the IBM FlashSystem A9000 and A9000R Open API Common Information Model (CIM) agent does not have any features that use this method.

Parameters

Table 1 describes the parameters of the SetProperty method.

Table 1. SetProperty method parameters

Parameter	Type	Description
InstanceName	String	Defines the name of the instance.
Property	String	The name of the property whose value is to be returned from the instance.

Return values

If successful, the instance is updated. Otherwise, one of the following return codes is returned:

- CIM_ERR_ACCESS_DENIED
- CIM_ERR_INVALID_NAMESPACE
- CIM_ERR_INVALID_PARAMETER
- CIM_ERR_INVALID_CLASS
- CIM_ERR_NOT_FOUND
- CIM_ERR_NO_SUCH_PROPERTY
- CIM_ERR_TYPE_MISMATCH
- CIM_ERR_FAILED

Retrieving a single qualifier declaration from the target namespace

The GetQualifier method retrieves a single qualifier declaration from the target namespace.

Parameters

Table 1 describes the parameters of the GetQualifier method.

Table 1. GetQualifier method parameters

Parameter	Type	Description
QualifierName	String	Defines the qualifier whose declaration is to be returned.

Return values

If successful, the value of the qualifier is returned. Otherwise, one of the following return codes is returned:

- CIM_ERR_ACCESS_DENIED
- CIM_ERR_INVALID_NAMESPACE
- CIM_ERR_INVALID_PARAMETER
- CIM_ERR_NOT_FOUND
- CIM_ERR_FAILED

Creating or modifying a qualifier declaration in the target namespace

The SetQualifier method creates or updates a qualifier declaration in the target namespace.

Parameters

Table 1 describes the parameters of the SetQualifier method.

Table 1. SetQualifier method parameters

Parameter	Type	Description
QualifierDeclaration	Void	Defines the qualifier declaration to be added to the target namespace.

Return values

If successful, the qualifier is updated in the target namespace. Otherwise, one of the following error codes is returned:

- CIM_ERR_ACCESS_DENIED
- CIM_ERR_INVALID_NAMESPACE
- CIM_ERR_INVALID_PARAMETER
- CIM_ERR_NOT_FOUND
- CIM_ERR_FAILED

Enumerating qualifier declarations from the target namespace

The EnumerateQualifiers method enumerates qualifier declarations from the target namespace.

There are no parameters for this method.

Return values

If successful, zero or more qualifier declarations are returned. Otherwise, one of the following error codes is returned:

- CIM_ERR_ACCESS_DENIED
- CIM_ERR_INVALID_NAMESPACE
- CIM_ERR_INVALID_PARAMETER
- CIM_ERR_FAILED

Common Information Model (CIM) agent communication methods that cannot be used

The following Common Information Model (CIM) agent communication methods cannot be used because they are not supported by IBM FlashSystem® A9000 and A9000R:

- **DeleteClass method (not supported)**
  The DeleteClass method deletes a single class from the target namespace, but it is not supported and cannot be used. The CIM_ERR_NOT_SUPPORTED error code is returned to the client application if a request to process this operation is received.
- **CreateClass method (not supported)**
  The CreateClass method creates a class from the target namespace, but it is not supported and cannot be used. The CIM_ERR_NOT_SUPPORTED error code is returned to the client application if a request to process this operation is received.
- **ModifyClass method (not supported)**
  The ModifyClass method modifies an existing class, but it is not supported and cannot be used. The CIM_ERR_NOT_SUPPORTED error code is returned to the client application if a request to process this operation is received.
- **DeleteQualifier method (not supported)**
  The DeleteQualifier method deletes a single class from the target namespace, but it is not supported and cannot be used. The CIM_ERR_NOT_SUPPORTED error code is returned to the client application if a request to process this operation is received.

DeleteClass method (not supported)

The DeleteClass method deletes a single class from the target namespace, but it is not supported and cannot be used. The CIM_ERR_NOT_SUPPORTED error code is returned to the client application if a request to process this operation is received.

CreateClass method (not supported)

The CreateClass method creates a class from the target namespace, but it is not supported and cannot be used. The CIM_ERR_NOT_SUPPORTED error code is returned to the client application if a request to process this operation is received.

ModifyClass method (not supported)

The ModifyClass method modifies an existing class, but it is not supported and cannot be used. The CIM_ERR_NOT_SUPPORTED error code is returned to the client application if a request to process this operation is received.

DeleteQualifier method (not supported)

The DeleteQualifier method deletes a single class from the target namespace, but it is not supported and cannot be used. The CIM_ERR_NOT_SUPPORTED error code is returned to the client application if a request to process this operation is received.

Return error codes

The Common Information Model object manager (CIMOM) returns status to the client application.

The return status is sent to the client application in one of the following ways:

- Through HTTP status messages
- Through error codes that are contained within `<METHODRESPONSE>` or `<IMETHODRESPONSE>` XML tags.

Table 1 describes the vendor-specific status codes that the CIMOM might return. For CIM standard return codes, see the CIM schema.

Table 1. Return error codes for the CIMOM

Code	Symbolic Name	Definition
1	CIM_ERR_FAILED	A general error occurred that is not covered by a more specific error code.
2	CIM_ERR_ACCESS_DENIED	Access to a CIM resource was not available to the client.
3	CIM_ERR_INVALID_NAMESPACE	The target namespace does not exist.
4	CIM_ERR_INVALID_PARAMETER	One or more parameter values that are passed to the method were not valid.
5	CIM_ERR_INVALID_CLASS	The specified class does not exist.
6	CIM_ERR_NOT_FOUND	The requested object was not found.
7	CIM_ERR_NOT_SUPPORTED	The requested operation is not supported.
8	CIM_ERR_CLASS_HAS_CHILDREN	The operation cannot be carried out on this class because it has subclasses.
9	CIM_ERR_CLASS_HAS_INSTANCES	The operation cannot be carried out on this class because it has instances.
10	CIM_ERR_INVALID_SUPERCLASS	The operation cannot be carried out because the specified superclass does not exist.
11	CIM_ERR_ALREADY_EXISTS	The operation cannot be carried out because the object exists.
12	CIM_ERR_NO_SUCH_PROPERTY	The specified property does not exist.
13	CIM_ERR_TYPE_MISMATCH	The value that is supplied is not compatible with the type that is specified.
14	CIM_ERR_QUERY_LANGUAGE_NOT_SUPPORTED	The query language is not recognized or supported.
15	CIM_ERR_INVALID_QUERY	The query is not valid for the specified query language.
16	CIM_ERR_METHOD_NOT_AVAILABLE	The extrinsic method cannot be run.
17	CIM_ERR_METHOD_NOT_FOUND	The specified extrinsic method does not exist.
20	CIM_ERR_LOW_ON_MEMORY	There is not enough memory.
21	XMLERROR	An XML error occurred.
22	CIM_ERR_LISTNER_ALREADY_DEFINED	The listener is already defined.
23	CIM_ERR_INDICATION_NOT_COLLECTED	The indications are not collected.
24	CIM_ERR_NO_METHOD_NAME	The method name is null.
25	CIM_ERR_INVALID_QUALIFIER_DATATYPE	The data type qualifier is not valid.
26	CIM_ERR_NAMESPACE_NOT_IN_MANAGER	The namespace value is not found.
27	CIM_ERR_INSTANTIATE_FAILED	The instantiation failed.
28	CIM_ERR_FAILED_TO_LOCATE_INDICATION_ HANDLER	The indication handler is not found.
29	CIM_ERR_IO_EXCEPTION	An I/O exception occurred.
30	CIM_ERR_COULD_NOT_DELETE_FILE	The file cannot be deleted.
31	INVALID_QUALIFIER_NAME	The qualifier name is null.
32	NO_QUALIFIER_VALUE	The qualifier value is null.
33	NO_SUCH_QUALIFIER1	There is no such qualifier.
34	NO_SUCH_QUALIFIER2	There is no such qualifier.
35	QUALIFIER_UNOVERRIDABLE	The qualifier cannot be overwritten.
36	SCOPE_ERROR	A scope error occurred.
37	TYPE_ERROR	A type error occurred.
38	CIM_ERR_MISSING_KEY	The key is missing.
39	CIM_ERR_KEY_CANNOT_MODIFY	The key cannot be modified.
40	CIM_ERR_NO_KEYS	There are no keys found.
41	CIM_ERR_KEYS_NOT_UNIQUE	The keys are not unique.
100	CIM_ERR_SET_CLASS_NOT_SUPPORTED	The set class operation is not supported.
101	CIM_ERR_SET_INSTANCE_NOT_SUPPORTED	The set instance operation is not supported.
102	CIM_ERR_QUALIFIER_NOT_FOUND	The qualifier value is not found.
103	CIM_ERR_QUALIFIERTYPE_NOT_FOUND	The qualifier type is not found.
104	CIM_ERR_CONNECTION_FAILURE	The connection failed.

Code	Symbolic Name	Definition
105	CIM_ERR_FAIL_TO_WRITE_TO_SERVER	There is a fail to write to the server.
106	CIM_ERR_SERVER_NOT_SPECIFIED	The server is not specified.
107	CIM_ERR_INDICATION_ERROR	There is an indication processing error.
108	CIM_ERR_FAIL_TO_WRITE_TO_CIMOM	A write operation to the CIMOM failed.
109	CIM_ERR_SUBSCRIPTION_EXISTS	A subscription exists.
110	CIM_ERR_INVALID_SUBSCRIPTION_DEST	The subscription destination is not valid.
111	CIM_ERR_INVALID_FILTER_PATH	The filter path is not valid.
112	CIM_ERR_INVALID_HANDLER_PATH	The handler path is not valid.
113	CIM_ERR_NO_FILTER_INSTANCE	The filter instance is not found.
114	CIM_ERR_NO_HANDLER_INSTANCE	The handler instance is not found.
115	CIM_ERR_UNSUPPPORTED_FILTER	The filter that is referenced in the subscription is not supported.
116	CIM_ERR_INVALID_TRUSTSTORE	The CIMOM cannot be connected to because there is a bad or missing truststore or an incorrect truststore password.
117	CIM_ERR_ALREADY_CONNECTED	The CIMOM cannot be connected to because it is already connected.
118	CIM_ERR_UNKNOWN_SERVER	The server is unknown. The CIMOM cannot accept connections.
119	CIM_ERR_INVALID_CERTIFICATE	The correct certificate cannot be found in the truststore. The CIMOM cannot accept connections.

Functional profiles, diagrams, and methods

The following sections detail the functional profiles, diagrams, and methods of the Common Information Model (CIM).

The functional diagrams show specific functions that the CIM agent provides and illustrate the architecture of the CIM agent for IBM FlashSystem® A9000 and A9000R systems. The method sections describe the uses of and parameters associated with the methods associated with each profile and matching diagram(s).

- **Block Server performance profile**
  Block server performance is the Storage Management Initiative Specification (SMI-S) subprofile that describes how to present performance statistics. It uses a number of metrics with definitions that are standardized for all storage systems that are represented by SMI-S. These metrics are defined in the CIM_BlockStorageStatisticalData class definition.
- **Block Services profile**
  The IBM FlashSystem A9000 and A9000R system architecture provides a means for you to customize the underlying resources.
- **iSCSI Target Ports profile**
  The iSCSI Target Ports subprofile provides a standard model for representing iSCSI elements as Common Information Model (CIM) objects.
- **Masking and Mapping profile**
  You can control which client hosts can see and access storage volumes by mapping and masking client hosts. The Storage Networking Industry Association (SNIA) Storage Management Initiative Specification (SMI-S) specifications define the Mapping and Masking profile to support this function.
- **Indication profile**
  The Storage Management Initiative Specification (SMI-S) supports two types of indications: Lifecycle indications and alert indications.
- **Replication Services profile**
  The Replication Services profile allows a storage system to copy data from a source element to a target element. The copy operations can be initiated on elements from the same storage system or across a connection to a different storage system.
- **Job Control profile**
  In some profiles, like the replication services profile, some of the methods described may take some time to execute. In these cases, a mechanism is needed to handle asynchronous execution of the method as a 'job.' The Job Control profile defines the constructs and behavior for job control.
- **Thin Provisioning profile**
  Thin provisioning is a capability of some Block Services implementations. It defers provisioning of backing store for regions of a volume until the regions are accessed (written) by the consumer (for example, host file system).

Block Server performance profile

Block server performance is the Storage Management Initiative Specification (SMI-S) subprofile that describes how to present performance statistics. It uses a number of metrics with definitions that are standardized for all storage systems that are represented by SMI-S. These metrics are defined in the CIM_BlockStorageStatisticalData class definition.

For the Common Information Model (CIM) agent for IBM FlashSystem® A9000 and A9000R systems, statistics are only provided for the volumes and hosts.

Table 1 provides details about the specific metrics that are supported by IBM FlashSystem A9000 and A9000R arrays and their components. In addition, the table provides a list of ElementType values that correspond to the components of IBM FlashSystem A9000 and A9000R arrays.

Table 1. Block Server Performance metrics

	Equivalent CIM class	Corresponding ElementType value	Properties that are supplied in associated CIM_BlockStorage StatisticalData instances

	Equivalent CIM class	Corresponding ElementType value	Properties that are supplied in associated CIM_BlockStorage StatisticalData instances
Volume	IBMTSDS_ SEVolumeStatistics	8	TotalIOs KBytesTransferred KBytesRead KBytesWritten ReadIOs ReadHitIOs WriteIOs WriteHitIOs IOTimeCounter ReadIOTimeCounter WriteIOTimeCounter ReadHitIOTimeCounter WriteHitIOTimeCounter
Host	IBMTSDS_HostStatistics	4	TotalIOs KBytesTransferred KBytesRead KBytesWritten ReadIOs ReadHitIOs WriteIOs WriteHitIOs IOTimeCounter ReadIOTimeCounter WriteIOTimeCounter ReadHitIOTimeCounter WriteHitIOTimeCounter

- **Block Server Performance object model**
  Two major categories of classes in the object model determine the way clients retrieve performance statistics.
- **Block Server Performance methods**
  The following sections describe functional methods associated with the Block Server Performance profile and object model, including uses of and parameters used by each method.

Block Server Performance object model

Two major categories of classes in the object model determine the way clients retrieve performance statistics.

The following categories are used:

- The first category is a set of classes where each instance of the class represents a single performance statistics record. For example, the statistics for a single volume.
- The second is a set of classes that are required to use an extrinsic method to retrieve a string that contains a batch of performance statistics.

See Figure 1 for the Block Server Performance Storage Management Initiative Specification (SMI-S) model for IBM FlashSystem A9000 and A9000R systems.

Figure 1. Block Server Performance SMI-S model for IBM FlashSystem A9000 and A9000R systems

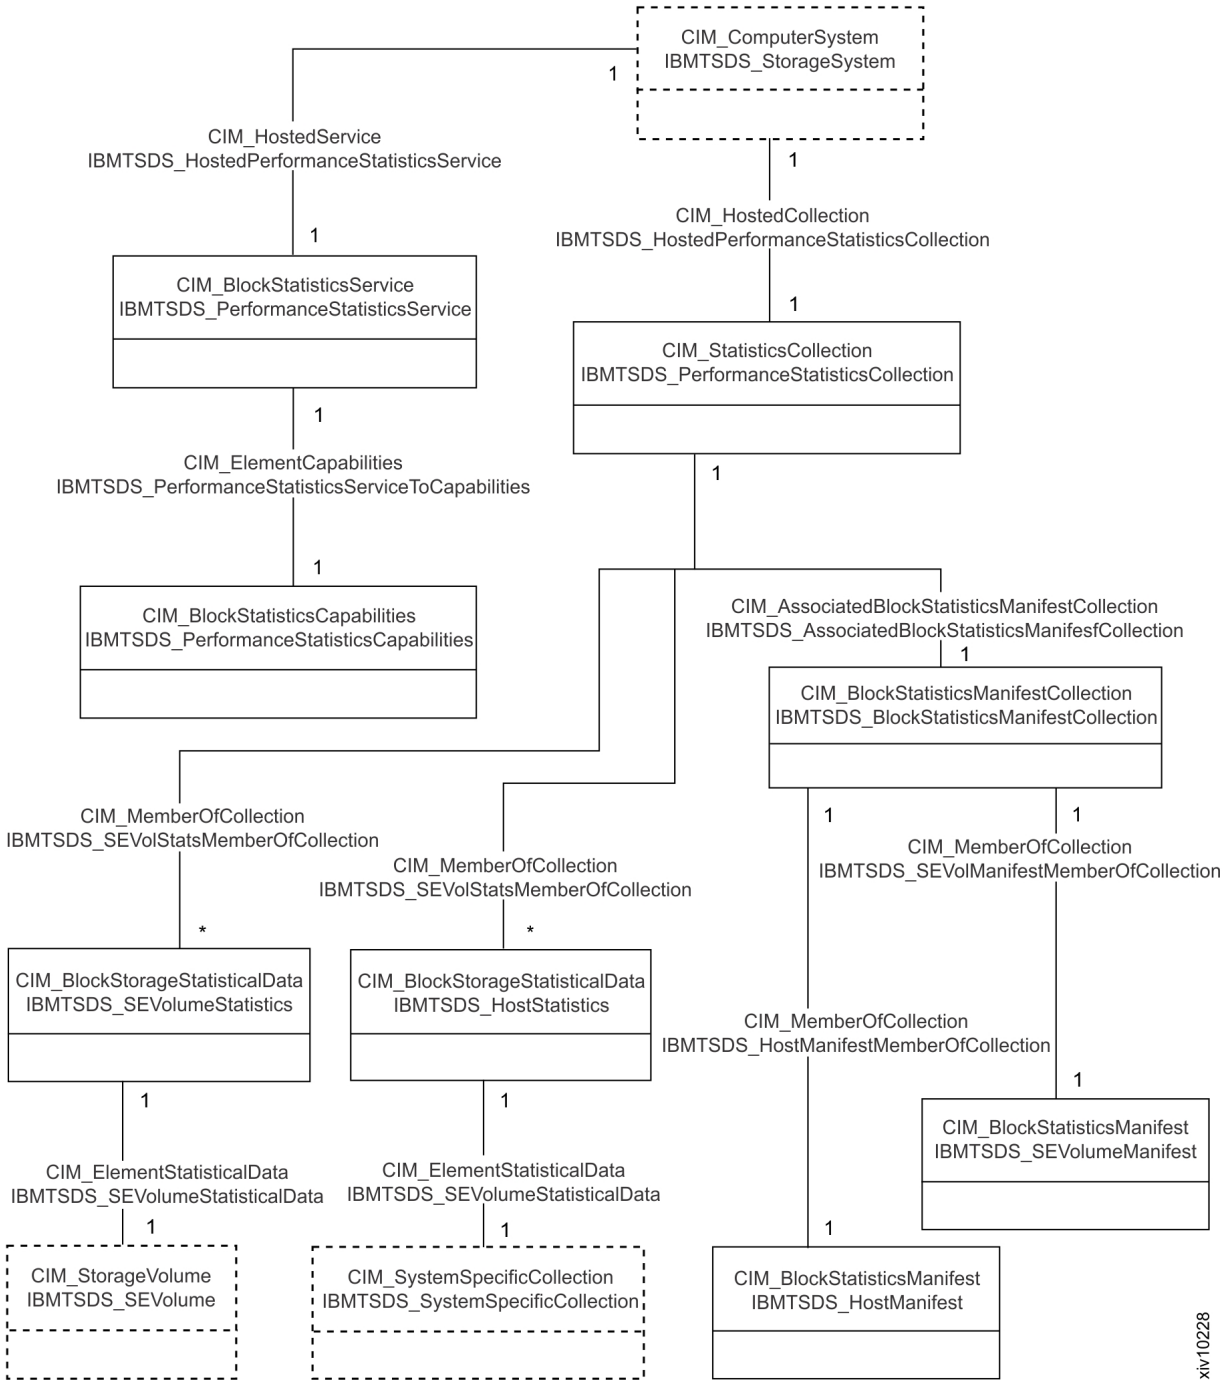

The classes that inherit from CIM_BlockStorageStatisticalData IBMTSDS_SEVolumeStatistics and IBMTSDS_HostStatistics are of the first category. Each instance contains properties that describe the performance statistics of a single instance of IBMTSDS_SEVolume or IBMTSDS_SystemSpecificCollection.

IBMTSDS_PerformanceStatisticsCollection does not fall under either of the categories, but is a class that associates all of the Block Server Performance-related classes for a particular system.

The rest of the classes are under the second category. They are used for batch gathering of statistics. IBMTSDS_PerformanceStatisticsService contains the GetStatisticsCollection method, which you use to get a string representation of a batch of statistics instances. IBMTSDS_PerformanceStatisticsCapabilities defines what the performance statistics-related capabilities of the Common Information Model (CIM) agent are.

The classes that inherit from CIM_BlockStatisticsManifest define filters for the statistics that are returned by GetStatisticsCollection.

IBMTSDS_BlockStatisticsManifestCollection is a group of BlockStatisticsManifest instances that gets passed into the GetStatisticsCollection method. The BlockStatisticsManifest classes contain Boolean properties, one for each property in the corresponding BlockStorageStatisticalData class.

If a BlockStorageStatisticalData had a statistics property XXX, then the corresponding BlockStatisticsManifest class would have a boolean property called IncludeXXX. If IncludeXXX was set to true, GetStatisticsCollection would return the data for XXX. If IncludeXXX was set to false, GetStatisticsCollection would not return the data for XXX.

The IBM FlashSystem A9000 and A9000R CIM agent does not allow clients to pick and choose which properties it is interested in. Therefore, all of the IncludeXXX properties for the XXX properties that the CIM agent supports are set to true.

All statistics attributes are in units of kilobytes, milliseconds, or just a count. For example, the number of I/O operations. However, the statistics are just running counters. They do not provide information about rates. For example, I/O operations per second. When the counters reach an internal limit, they roll back to zero. If the client application is monitoring these statistics at a constant rate to calculate I/O rates, it must know when the counters roll back to zero.

Block Server Performance methods

The following sections describe functional methods associated with the Block Server Performance profile and object model, including uses of and parameters used by each method.

- Obtaining performance statistics data
- Obtaining volume or host statistics

- **Obtaining performance statistics data**
  The IBMTSDS_PerformanceStatisticsService.GetStatisticsCollection method returns a string representation of a set of performance statistics data.
- **Obtaining volume or host statistics**
  You can use the IBMTSDS_PerformanceStatisticsService.GetStatisticsCollection method to obtain volume or host statistics.

Obtaining performance statistics data

The IBMTSDS_PerformanceStatisticsService.GetStatisticsCollection method returns a string representation of a set of performance statistics data.

Parameters

The following list describes the parameters of the IBMTSDS_PerformanceStatisticsService.GetStatisticsCollection method.

ElementTypes
: An array of values that indicates the type of element that returns statistics. Volumes and hosts use the standard value in the Common Information Model (CIM) schema. If this parameter is left null, the default value 4 and 8 are used. See the **`BlockServerPerformance.mof`** file for details.

ManifestCollection *(required)*
: A reference to an instance of the collection that represents a collection of BlockStatisticsManifest instances to use to filter the output. This parameter is not a reference to BlockStatisticsManifest. It is a ManifestCollection because the method can return statistics for multiple ElementTypes values at the same time, and each ElementType instance has its own BlockStatisticsManifest. The CIM agent currently supplies one instance of ManifestCollection, which is the default ManifestCollection.

Statistics
: An output parameter that is an array of strings that represents a batch of performance data. Each array element is an instance of a statistics class. Each array element is formatted as a semicolon-separated list of values. The order of the returned values matches the properties definition order in the corresponding CIM_BlockStatisticsManifest class in mof. For example, the statistics output parameter might display as follows:

```
IBM.2810-6000095-100916;8;20110118225313.890995+480;3111;166920;20864419;1426;
1422;1119849;1078871;67382;1685;1673;19744570;19692176;99538;
IBM.2810-6000095-100930;8;20110118225313.891038+480;8386;504537;76662400;1426;
1422;1182723;1133669;67382;6960;6940;75479677;75360317;437155;
```

where each row is an array element. The ElementType value of each row is 8, which refers to volumes. So there are two instances of VolumeStatistics with attributes of each in semicolon-separated strings, in the order they are defined in the IBMTSDS_SEVolumeManifest class in BlockServerPerformance.mof.

```
{
StatisticTime;TotalIOs;KBytesTransferred;IOTimeCounter;ReadIOs;ReadHitIOs;
ReadIOTimeCounter;ReadHitIOTimeCounter;KBytesRead;WriteIOs;WriteHitIOs;
WriteIOTimeCounter;WriteHitIOTimeCounter;KBytesWritten; }
```

Obtaining volume or host statistics

You can use the IBMTSDS_PerformanceStatisticsService.GetStatisticsCollection method to obtain volume or host statistics.

Procedure

1. Use enumerateInstanceNames on IBMTSDS_PerformanceStatisticsService and then save the reference.
2. Use enumerateInstancesNames on IBMTSDS_BlockStatisticsManifestCollection and then save the reference.
3. Use enumerateInstances on IBMTSDS_VolumeManifest (or CIM_BlockStatisticsManifest and look for the ElementType value 8) or IBMTSDS_SystemSpecificCollection (or CIM_BlockStatisticsManifest and look for the ElementType value 4), and then save BulkFormat.
4. Use invokeMethod on GetStatisticsCollection with the ElementType value set to 8 for volume or to 4 for host (if ElementType is null, ElementType values of 4 and 8 are assigned), the StatisticsFormat value set to 2, the ManifestCollection value obtained in the second step, and the PerformanceStatisticsCollection instance value that is obtained in the first step.
5. Check the output statistics (Statistics). For each array element in the statistics output, string tokenize on semicolon, and save the properties in a table. Check the properties.

Block Services profile

The IBM FlashSystem® A9000 and A9000R system architecture provides a means for you to customize the underlying resources.

The following definitions describe the various layers of abstraction that make this customization possible:

The primordial storage pool
: Logical entity that contains all the available unformatted or unprepared disk capacity on an array. Device Concrete StoragePool instances are allocated from Primordial StoragePool. IBM FlashSystem A9000 and A9000R arrays have only one Primordial StoragePool instance per array.

Concrete storage pools
: Logical entities that are allocated from Primordial StoragePool. Concrete StoragePool instances enable storage administrators to manage relationships between volumes and snapshots and to define separate capacity provisioning and snapshot requirements for separate applications and departments. Storage pools are not tied to specific physical resources, nor are they part of the data-distribution scheme.
IBM FlashSystem A9000 and A9000R support two types of concrete storage pools: VirtualPool and SnapshotPool. A VirtualPool instance is allocated from Primordial StoragePool directly. It contains volumes and at most one SnapshotPool instance.

  A SnapshotPool instance is allocated from a VirtualPool instance and contains snapshots of volumes in the VirtualPool instance in which this SnapshotPool instance is located. A VirtualPool instance without its own SnapshotPool instance cannot contain snapshots

  .

Storage volumes
: Logical units that can be mapped to a host or a cluster. Volumes are created from storage pools and are managed within the context of storage pools. For the IBM FlashSystem A9000 and A9000R systems, volumes do not occupy physical capacity when created. Space is allocated when data is written to the volume. Moreover, volumes are not exclusively associated with a subset of physical resources, nor is there a permanent static relationship between logical volumes and specific physical resources.

See *IBM FlashSystem A9000 and A9000R: Architecture, Implementation, and Usage* on the IBM® Redbooks® website and the *IBM FlashSystem A9000 Product Overview* or *IBM FlashSystem A9000R Product Overview* for more information about the architecture of IBM FlashSystem A9000 and A9000R systems.

- **Block Services object model**
  The Common Information Model (CIM) agent is designed to allow for the retrieval of information for each layer of abstraction.
- **Block Services methods**
  The following sections describe functional methods associated with the Block Services profile and object model, including uses of and parameters used by each method.

Block Services object model

The Common Information Model (CIM) agent is designed to allow for the retrieval of information for each layer of abstraction.

Starting with the 12.0.1 release, the Storage Management Initiative Specification (SMI-S) specification 1.6 is fully supported. You can create, delete, or modify services.

See Figure 1 for the Block Services SMI-S model for IBM FlashSystem A9000 and A9000R systems.

Figure 1. Block Services SMI-S model for IBM FlashSystem A9000 and A9000R systems

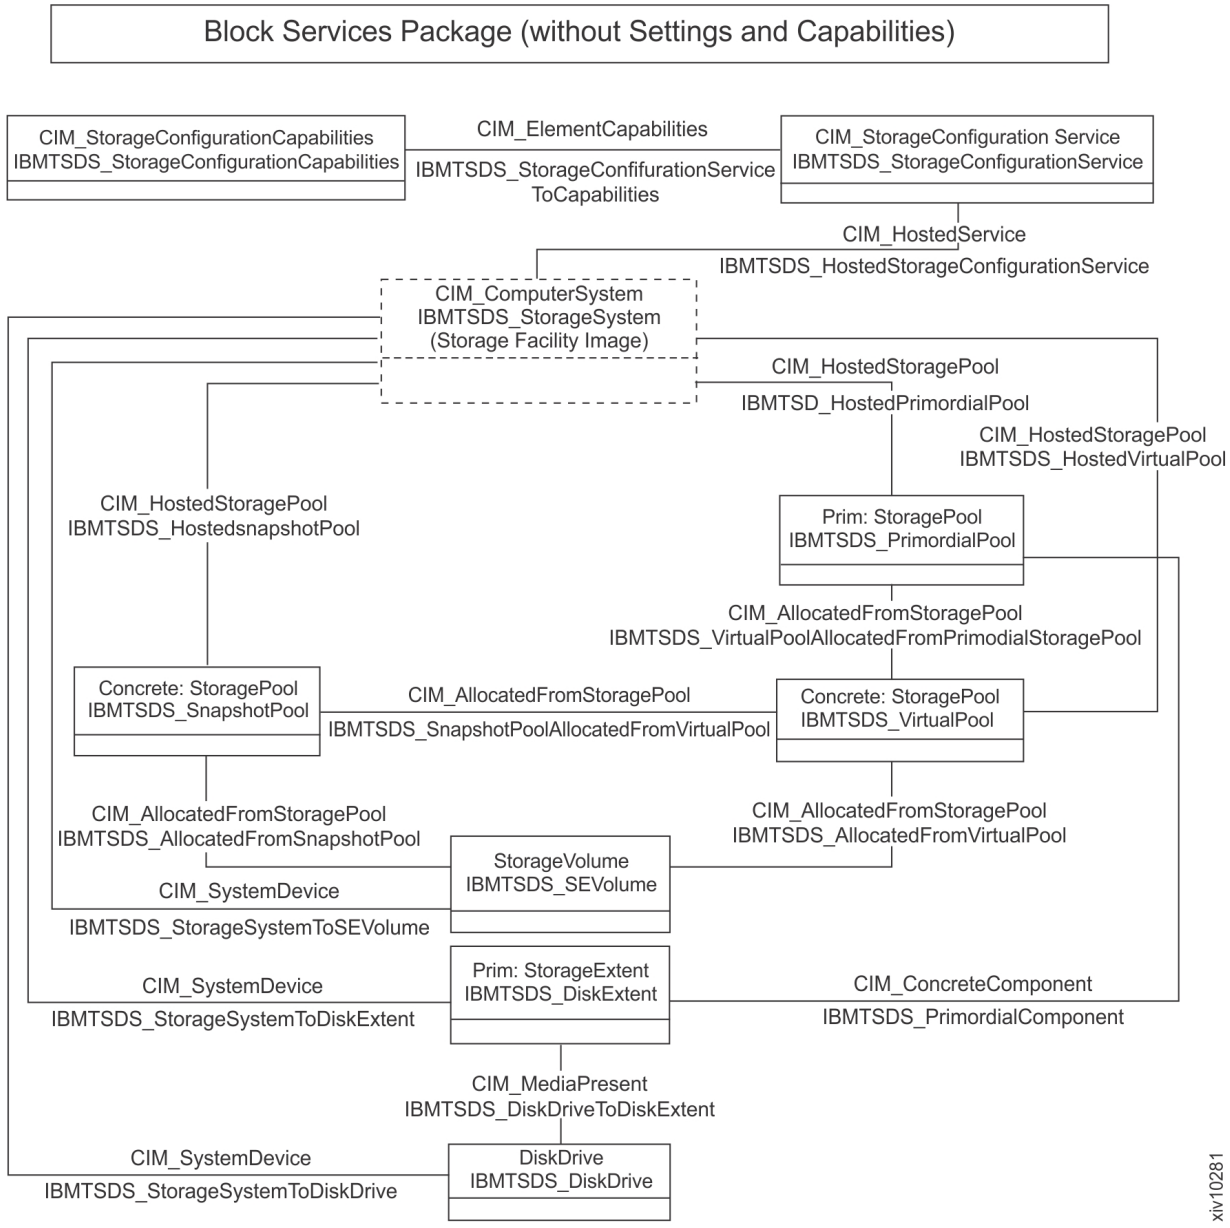

As required by SMI-S, each instance of storage pool has a corresponding instance of Storage Capabilities. Each instance of Storage Capabilities is associated with the instances of Storage Setting that are valid for creating volumes from the associated storage pool. See Figure 2 for the Block Services package with settings and capabilities model.

Figure 2. Block Services Package with Settings and Capabilities model

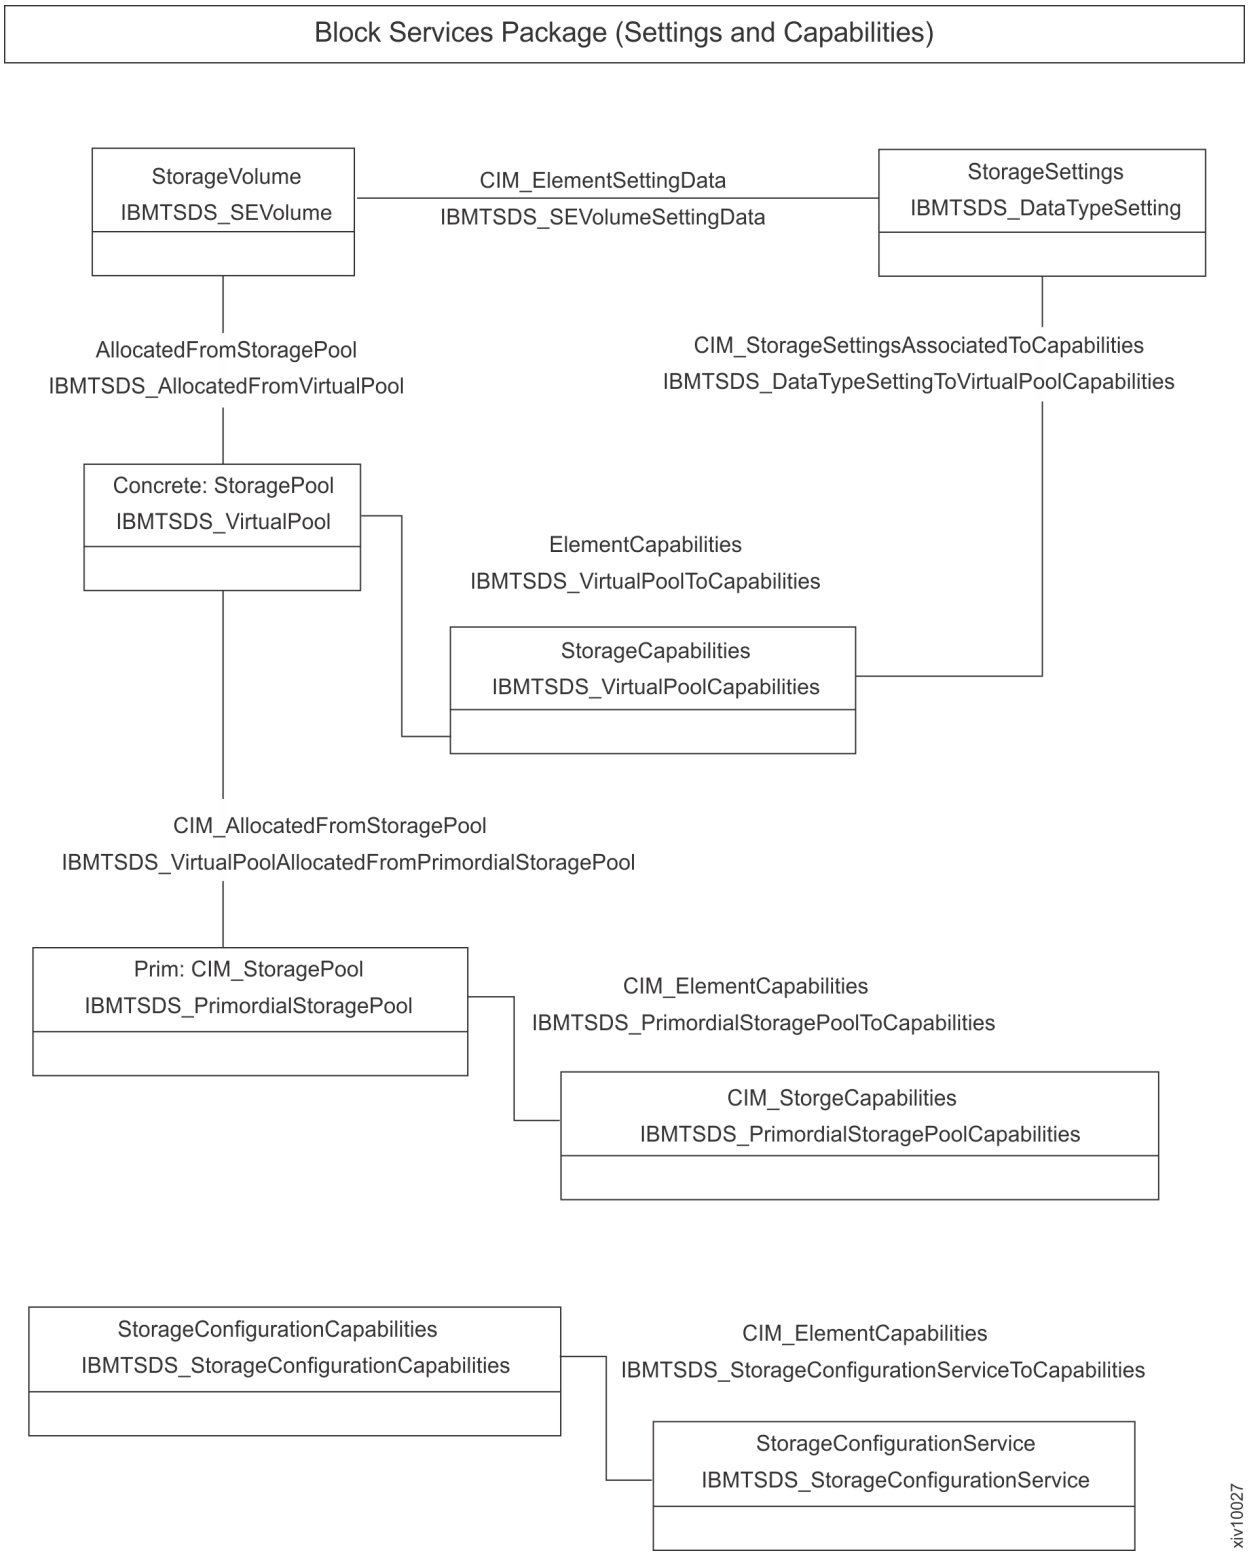

StorageSetting is a class that contains properties to specify the quality of service, such as DataRedundancy or parity layout. When you create a StorageVolume or Concrete StoragePool instance, a StorageSetting instance is supplied as the Goal parameter for the appropriate method.

The IBM FlashSystem A9000 and A9000R CIM agent has four CIM_StorageSetting instances defined:

IBMTSDS_VirtualPoolSetting.InstanceID="IBMTSDS:IBM XIV Virtual Storage Pool Setting"
This StorageSetting instance is supplied as the Goal parameter for the CreateOrModifyStoragePool method to create or modify a VirtualPool instance.
IBMTSDS_SnapshotPoolSetting.InstanceID="IBMTSDS:IBM XIV Snapshot Pool Setting"
This StorageSetting instance is supplied as the Goal parameter for the CreateOrModifyStoragePool method to create or modify a SnapshotPool instance.
IBMTSDS_DataTypeSetting.InstanceID="IBMTSDS:XIVBlockSize"
This StorageSetting instance is supplied as the Goal parameter for the CreateOrModifyElementFromStoragePool and CreateOrModifyElementsFromStoragePool methods to create or modify StorageVolume instances with size in terms of GB.
The system allocates the soft volume size as the minimum number of discrete 1 GB increments needed to meet the requested volume size.

IBMTSDS_DataTypeSetting.InstanceID="IBMTSDS:SystemBlockSize"

This StorageSetting instance is supplied as the Goal parameter for the CreateOrModifyElementFromStoragePool and CreateOrModifyElementsFromStoragePool methods to create or modify StorageVolume instances with size in terms of blocks.
The volume's capacity is indicated as a discrete number of 512-byte blocks. The system allocates the soft volume size that is used within the StoragePool instance as the minimum number of discrete 1 GB increments to meet the requested size. However, the size that is reported to hosts is equivalent to the precise number of blocks defined.

For more information about the block services object model, see the BlockServices.mof file.

Extrinsic Methods

Use an associator's CIM client request and the association class CIM_ElementCapabilities between an instance of CIM_StorageConfigurationService and the result class CIM_StorageConfigurationCapabilities.

The SupportedSynchronousActions and SupportedAsynchronousActions properties identify whether the action is initiated synchronously or asynchronously. The IBM FlashSystem A9000 and A9000R CIM agent supports the following synchronous actions in the IBMTSDS_StorageConfigurationService class:

Table 1. Synchronous actions

Synchronous actions	IBM FlashSystem A9000 and A9000R CIM functions
StoragePool Creation	CreateOrModifyStoragePool
StoragePool Modification	CreateOrModifyStoragePool
StoragePool Deletion	DeleteStoragePool
StorageVolume Creation	CreateOrModifyElementFromStoragePool CreateorModifyElementsFromStoragePool
StorageVolume Modification	CreateOrModifyElementFromStoragePool CreateorModifyElementsFromStoragePool
StorageVolume Expansion	CreateOrModifyElementFromStoragePool CreateorModifyElementsFromStoragePool
StorageVolume Shrinking	CreateOrModifyElementFromStoragePool CreateorModifyElementsFromStoragePool
StorageVolume Deletion	Return to StoragePool
StorageVolume Deletion	Return Elements to StoragePool

The following functions can be used to determine what sizes of Concrete StoragePool and StorageVolume instances can be created:

IBMTSDS_PrimordialStoragePool.GetSupportedSizeRange
: Determines the supported physical size of Concrete StoragePool instances that can be created in Primordial StoragePool.

IBMTSDS_VirtualPool.GetSupportedSizeRange
: Determines the supported soft size of StorageVolume instances that can be created in a Concrete StoragePool instance.

Block Services methods

The following sections describe functional methods associated with the Block Services profile and object model, including uses of and parameters used by each method.

- **Creating or modifying a virtual pool or snapshot pool**
  The IBMTSDS_StorageConfigurationService. CreateOrModifyStoragePool method creates or modifies a VirtualPool instance or a SnapshotPool instance. The Primordial StoragePool instance cannot be created or modified.
- **Deleting a storage pool**
  The IBMTSDS_StorageConfigurationService. DeleteStoragePool method is used to delete a single storage pool. This method requires a single parameter. Specify the reference to the Concrete StoragePool instance that is to be deleted. Only a single StoragePool instance can be deleted at a time.
- **Creating, modifying, and/or moving multiple volumes in a single method call**
  IBMTSDS_StorageConfigurationService. CreateOrModifyElementsFromStoragePool is a vendor-extension method that allows for the creation / modification / moving of multiple volumes in a single method call.
- **Creating, modifying, and/or moving a single volume**
  IBMTSDS_StorageConfigurationService. CreateOrModifyElementFromStoragePool is the Common Information Model (CIM) and Storage Management Initiative Specification (SMI-S) standard method for creating or modifying a StorageVolume instance from a Concrete StoragePool instance.
- **Deleting a single volume**
  The IBMTSDS_StorageConfigurationService. ReturnToStoragePool method is used to delete a single volume. The only input parameter that is required is a reference to the volume that you want to delete. Only one volume can be deleted at a time.
- **Deleting multiple volumes concurrently**
  The IBMTSDS_StorageConfigurationService. ReturnElementsToStoragePool method is used to delete multiple volumes concurrently. You can pass in an array of references of volumes to delete.
- **Determining sizes to use to create elements from the primordial storage pool**
  The IBMTSDS_PrimordialStoragePool. GetSupportedSizeRange method is used to find out what sizes to use to create elements from the primordial storage pool.
- **Determining sizes to use to create elements from the virtual storage pool**
  The IBMTSDS_VirtualPool.GetSupportedSizeRange method is used to find out what sizes to use to create elements from the virtual storage pool.
- **Creating a concrete storage pool**
  You can use the CreateOrModifyStoragePool method to create a Concrete StoragePool instance.
- **Creating a storage volume**
  You can use the CreateOrModifyElementFromStoragePool method to create a StorageVolume instance.

Creating or modifying a virtual pool or snapshot pool

The IBMTSDS_StorageConfigurationService. CreateOrModifyStoragePool method creates or modifies a VirtualPool instance or a SnapshotPool instance. The Primordial StoragePool instance cannot be created or modified.

Parameters

The following list describes the parameters of the IBMTSDS_StorageConfigurationService. CreateOrModifyStoragePool method.

Pool *(required for modification)*
: As an input parameter, specifies whether you want to create or modify a pool. If you specify a reference to a pool, it indicates that you want to modify the pool. If the parameter is null, it indicates that you want to create pool.

InPools
: InPools specifies from which pool to create a pool. To create a VirtualPool instance, only the object reference of IBMTSDS_PrimordialStoragePool is allowed. All VirtualPool instances are created in Primordial StoragePool. To create a SnapshotPool instance, only the object reference of IBMTSDS_VirtualPool is allowed because a SnapshotPool instance is created in a VirtualPool instance.

  Note: The Common Information Model (CIM) schema defines this input parameter to be an array of strings that represent CIM object paths (COPs), and not actual references to objects.

Goal
: The Goal parameter represents the StorageSetting instance of the pool to be created. To create a VirtualPool instance, only the object reference IBMTSDS_VirtualPoolSetting is allowed; to create a SnapshotPool instance, only the object reference IBMTSDS_SnapshotPoolSetting is allowed. The Goal parameter also specifies the snapshot size value of the pool. See the BlockServer.mof file for details.

ElementName
: The ElementName property provides a means for you to set a meaningful name for the pool to be created or modified. If specified, limit it to 63 characters which can include letters, digits, blanks, -, _, . and ~ characters. Blanks cannot be the beginning and ending characters. If not specified during pool creation, a random pool name is generated in the form of `pool<random integer>`.

  Note: The name of the pool must be unique in the system. It cannot be a name that is already assigned to one of the other pools in the system.

Size *(required for creation)*
: As an input parameter, Size specifies the requested size of the pool. `Null` is not allowed for pool creation. As an output parameter, Size specifies the size achieved.

Deleting a storage pool

The IBMTSDS_StorageConfigurationService. DeleteStoragePool method is used to delete a single storage pool. This method requires a single parameter. Specify the reference to the Concrete StoragePool instance that is to be deleted. Only a single StoragePool instance can be deleted at a time.

Note: You must delete all volumes in the current pool before you delete a pool, and the Primordial StoragePool instance cannot be deleted.

Creating, modifying, and/or moving multiple volumes in a single method call

IBMTSDS_StorageConfigurationService. CreateOrModifyElementsFromStoragePool is a vendor-extension method that allows for the creation / modification / moving of multiple volumes in a single method call.

For large numbers of volumes, this method can be more efficient than calling CreateOrModifyElementFromStoragePool several times in a loop. The InPool, Goal, and ElementType input parameters are the same as in CreateOrModifyElementFromStoragePool.

Parameters

The following definitions summarize the parameters that are different from CreateOrModifyElementFromStoragePool.

TheElements *(required for volume modification or volume moving)*
: If this value is not null, it indicates that you want to modify each of the volumes specified.

Quantity *(required for volume creation)*
: As input, represents the number of StorageVolume instances to be created. This parameter must be null when you modify StorageVolume instances. As output, represents the number of volumes that are created if successful. Or, if creation was not successful because of too many volumes and not enough capacity on the VirtualPool instance, it represents the number of volumes that can be created.

ElementNames
: An array of the element names to assign to the various volumes that are being created or modified. Volume names can be supplied in two different ways:

  - The first way supplies each volume name with one element of the value. The length must be equal to the value for Quantity when you create volumes. If not `NULL`, the length must be the same as the length specified in the TheElements parameter when you modify volumes. For example, to create two volumes 'testVolume_3' and 'testVolume_4', set the ElementNames parameter's value as `['testVolume_3','testVolume_4']` and leave the value of the FirstSuffix parameter as `null`.
  - The second way supplies the leading string of volumes names as the only one element. In this case, names of all volumes have the same format `[LeadingString]_[IncrementingNumber]`. `[LeadingString]` is the value of this parameter, and the first `[IncrementingNumber]` is the value of the FirstSuffix parameter. The incremental number is 1. To create two volumes 'testVolume_3' and 'testVolume_4' in the first way, set the ElementNames parameter's value as `['testVolume']` and set the value of the FirstSuffix parameter as `3`.

FirstSuffix
: Represents the starting suffix number of volume names when you use the second way described to create volumes or modify volume names. Users can supply the starting `[IncrementingNumber]` with this parameter. If it is null, the default value is 1. When you use the first way described to create volumes or modify volume names, the value must be null.

  Note: The FirstSuffix parameter is not in the Common Information Model (CIM) or Storage Management Initiative Specification (SMI-S) schema. This parameter is an IBM® extension.

ReturnCodes
: Each volume that is created or modified can have a different error that is associated with it. This output parameter is an array of each of the individual return codes for each attempt.

Creating, modifying, and/or moving a single volume

IBMTSDS_StorageConfigurationService. CreateOrModifyElementFromStoragePool is the Common Information Model (CIM) and Storage Management Initiative Specification (SMI-S) standard method for creating or modifying a StorageVolume instance from a Concrete StoragePool instance.

This method can create or modify one volume at a time. This method is also used for moving a volume from one pool to another when the InPool parameter is specified with reference of the target pool.

Parameters

The following definitions summarize how the IBMTSDS_StorageConfigurationService. CreateOrModifyElementFromStoragePoolmethod handles different values for the input parameters:

TheElement *(required for modification)*
: As an input parameter, if the value is null, it indicates that you want to create a volume. If the value is not null, it indicates that you are trying to modify the specified volume. As an output parameter, TheElement references the volume that was created or modified.

ElementType
: An enumeration that indicates what type of element is being created or modified. Only volume values are supported. So the value of ElementType is 2 (StorageVolume) or 5 (ThinlyProvisionedStorageVolume). See the BlockServer.mof file for details.

InPool *(required for creation and moving)*
: A reference to the StoragePool instance which the volume is to be created in or moved to. It must be an instance of IBMTSDS_VirtualPool.

Goal
: A reference to the StorageSetting instance representing the DataType value of the volume. It must be a reference to an instance of IBMTSDS_DataTypeSetting. IBMTSDS_DataTypeSetting has two instances: IBMTSDS:SystemBlockSize and IBMTSDS:XIVBlockSize. In volume creation, if this parameter is null, the default value IBMTSDS:XIVBlockSize is used. In volume modification, if the Size parameter is specified, but this parameter is null, the default value IBMTSDS:XIVBlockSize is used, and the volume is resized.

Size
: In creating or modifying a volume, this parameter specifies the wanted size (as input) and the size achieved (as an output parameter). In modifying the size of a volume, only size expansion is supported by default. If you want to shrink a volume, set the ForceShrink parameter as true.

ForceShrink
: In modifying a volume, this parameter specifies a volume is shrunk if the Size parameter is smaller than the actual size of the volume. If this parameter is null or set to false, shrinking a volume fails. To shrink a volume, set this parameter to true. The default value is false.

  Attention: Shrinking a volume can cause data loss.

  The ForceShrink parameter is not in the CIM or SMI-S schema. This parameter is an IBM® extension.

Locked
: In modifying a volume/snapshot, this parameter specifies the volume/snapshot locking status to be locked or unlocked. If this parameter is null, the volume/snapshot locking status is not modified. If this parameter is set to 0, the volume/snapshot is unlocked. If this parameter is set to 1, the volume/snapshot is locked. Other values are invalid. When a volume/snapshot is locked, and other properties are modified in the same method, the volume/snapshot is locked in the last step. When a volume/snapshot is unlocked and other properties are modified in the same method, the volume/snapshot is unlocked in the first step.

  Note: The Locked parameter is not in the CIM or SMI-S schema. This parameter is an IBM extension.

ElementName
: A descriptive name to assign to the volume.

  Note: The volume name must be unique in the system. It cannot be a name that is already assigned to one of the other volumes in the system.

Deleting a single volume

The IBMTSDS_StorageConfigurationService. ReturnToStoragePool method is used to delete a single volume. The only input parameter that is required is a reference to the volume that you want to delete. Only one volume can be deleted at a time.

Deleting multiple volumes concurrently

The IBMTSDS_StorageConfigurationService. ReturnElementsToStoragePool method is used to delete multiple volumes concurrently. You can pass in an array of references of volumes to delete.

There is a ReturnCodes output parameter for the errors that can be generated for each individual volume that is being deleted.

Note: The ReturnElementsToStoragePool method is not in the Common Information Model (CIM) or Storage Management Initiative Specification (SMI-S) schema. This method is an IBM® extension.

Determining sizes to use to create elements from the primordial storage pool

The IBMTSDS_PrimordialStoragePool. GetSupportedSizeRange method is used to find out what sizes to use to create elements from the primordial storage pool.

CIM_StoragePool includes the standard methods GetSupportedSizes and GetSupportedSizeRange to perform this task, but for IBMTSDS_PrimordialStoragePool, only GetSupportedSizeRange is supported. It can be used to retrieve the supported physical size of a virtual pool which can be created. The Goal input parameter represents the pool data type to use to calculate the possible valid sizes. It returns a range and a divisor of valid sizes.

Note: If there is no spare space to create a minimum pool, the returned minimum and maximum size of GetSupportedSizeRange for the storage pool is zero.

Determining sizes to use to create elements from the virtual storage pool

The IBMTSDS_VirtualPool.GetSupportedSizeRange method is used to find out what sizes to use to create elements from the virtual storage pool.

This method can be used to retrieve the supported size of virtual pool which can be created. The Goal input parameter represents the pool data type used to calculate the possible valid sizes. It returns a range and a divisor of valid sizes.

Note: If there is no spare space to create a minimum pool, the returned minimum and maximum size of GetSupportedSizeRange for the storage pool is zero.

Creating a concrete storage pool

You can use the CreateOrModifyStoragePool method to create a Concrete StoragePool instance.

Procedure

1. Call IBMTSDS_PrimordialStoragePool.GetSupportedSizeRange with ElementType set to StoragePool. If the return code is `0` (success), present these values to the user and allow the user to pick a size.
2. Use InvokeMethod on IBMTSDS_StorageConfigurationService.CreateOrModifyStoragePool and specify the ElementName, Size, and Goal input parameters.
3. Save the pool output parameter. It is the reference to the VirtualPool instance that was created.

Creating a storage volume

You can use the CreateOrModifyElementFromStoragePool method to create a StorageVolume instance.

Procedure

1. Use the associatorNames parameter from the Concrete StoragePool instance that was created to CIM_StorageCapabilities to get a matching instance of IBMTSDS_VirtualPoolCapabilities.
2. Use the associatorNames parameter from the IBMTSDS_VirtualPoolCapabilities instance to CIM_StorageSetting to get valid instances of IBMTSDS_DataTypeSetting and allow user to pick the data type to use to create the volume.
3. Call IBMTSDS_VirtualPool.GetSupportedSizeRange on the virtual pool that is created in Step 1 using the data type that the user picked as the Goal input parameter. If the return code is `0` (success), present these values to the user and allow the user to pick a size.
4. Use InvokeMethod on IBMTSDS_StorageConfigurationService. CreateOrModifyElementFromStoragePool. Use the IBMTSDS_VirtualPool instance created in Step 1, the IBMTSDS_DataTypeSetting value, Size value, and ElementName value specified as the InPools, Size, Goal, and ElementName input parameters.
5. Save the TheElement output parameter and display it to the user, informing the user know that is the instance of the volume that was created.

iSCSI Target Ports profile

The iSCSI Target Ports subprofile provides a standard model for representing iSCSI elements as Common Information Model (CIM) objects.

A client can obtain information about iSCSI capabilities and settings of an IBM FlashSystem A9000 and A9000R storage array. A client can also obtain information that is provided by iSCSI-capable arrays and the storage volumes that are exposed by those arrays on an iSCSI network:

- Nodes
- Network Portals
- SCSI Ports

iSCSI sessions and settings are not currently configurable through the IBM FlashSystem A9000 and A9000R Storage Management Initiative Specification (SMI-S) provider.

- **iSCSI Target Ports Object Model**
  The Common Information Model (CIM) agent is designed to allow for the retrieval of information for each layer of abstraction.

- **iSCSI Target Ports methods**
  The following sections describe functional methods associated with the iSCSI Target Ports profile and object model, including uses of and parameters used by each method.

iSCSI Target Ports Object Model

The Common Information Model (CIM) agent is designed to allow for the retrieval of information for each layer of abstraction.

See Figure 1 for the iSCSI Target Ports Storage Management Initiative Specification (SMI-S) model for IBM FlashSystem A9000 and A9000R systems.

Figure 1. iSCSI Target Ports SMI-S model for IBM FlashSystem A9000 and A9000R systems

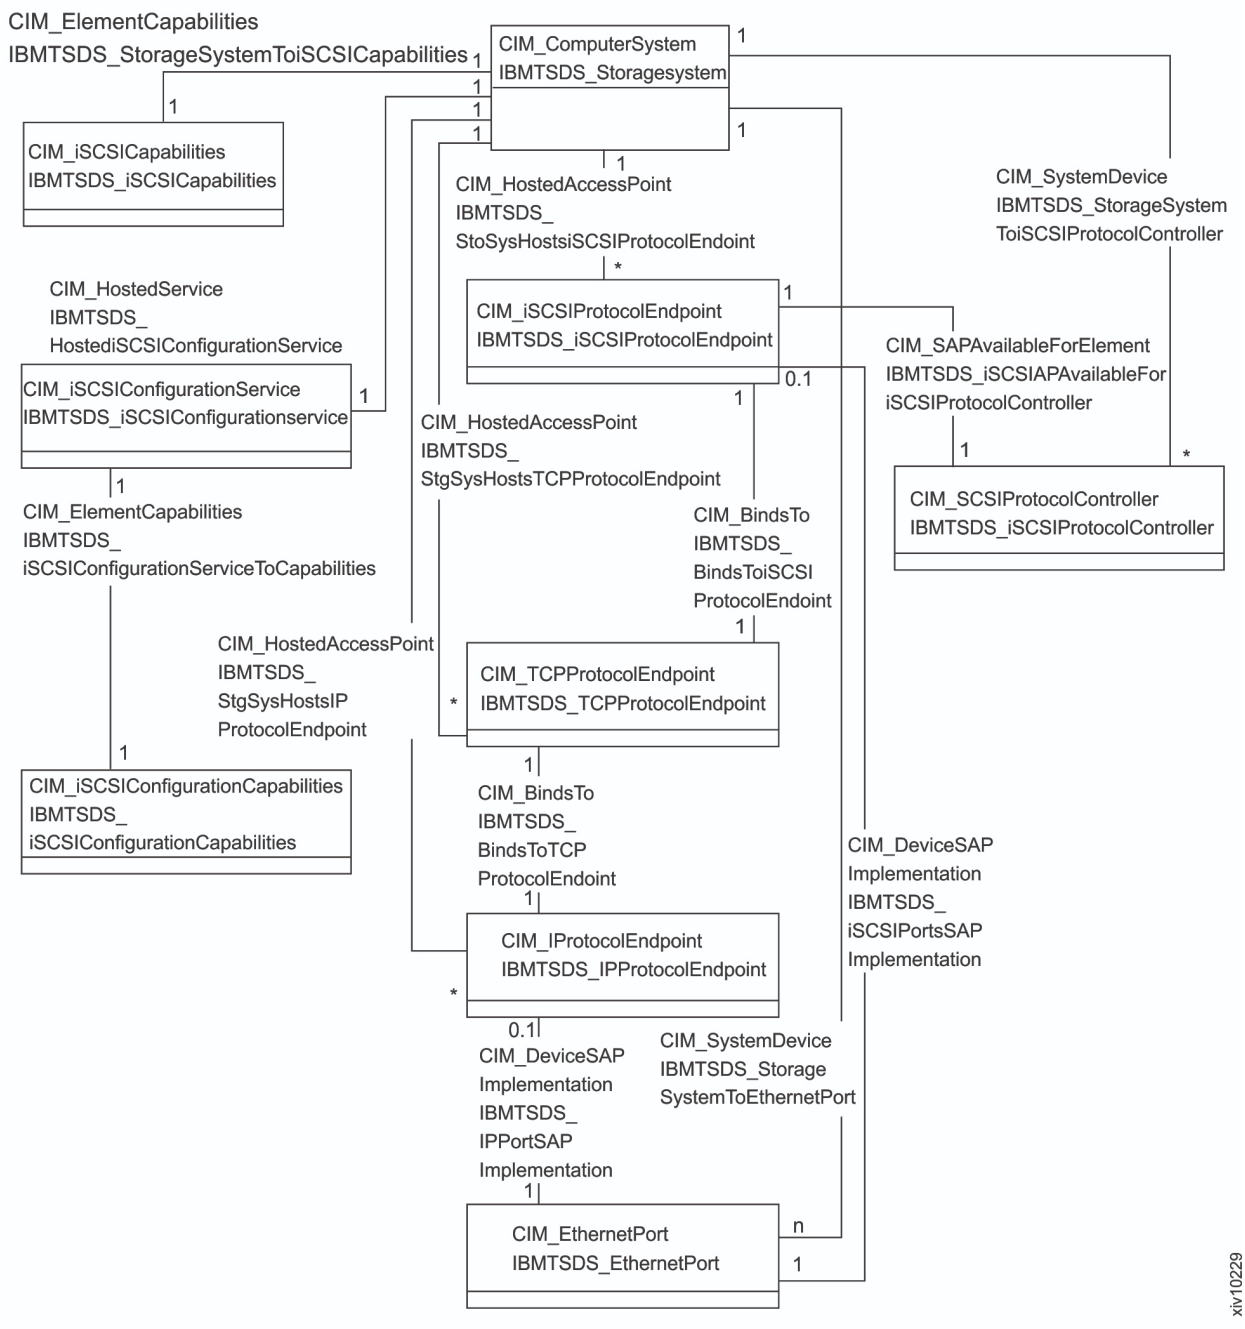

The following table provides a map of terminology from iSCSI standards and CIM class names on the previous UML diagram. iSCSI Session, Connection, and Portal Group are not supported by IBM FlashSystem A9000 and A9000R SMI-S.

Table 1. iSCSI terminology and CIM class names

iSCSI term	CIM class name	Description
Network entity	ComputerSystem	The Network Entity represents a device that is accessible from the IP network. A Network Entity has one or more Network Portals. Each Network Portal can be used to gain access to the IP network by some iSCSI Nodes that are contained in that Network Entity.

iSCSI term	CIM class name	Description
Node	SCSIProtocolController	The iSCSI Node represents a single iSCSI Target. There are one or more iSCSI Nodes within a Network Entity. The iSCSI Node is accessible through one or more Network Portals. An iSCSI Node is identified by its iSCSI Name. Separating the iSCSI Name from the addresses allows multiple iSCSI nodes to use the same address, and the same iSCSI node to use multiple addresses.
SCSI port	iSCSIProtocolEndpoint	A SCSI Port with an iSCSI service delivery subsystem. A collection of Network Portals that together acts as a SCSI Target or target.
Network Portal	TCPProtocolEndpoint, IPProtocolEndpoint, EthernetPort	A component of a Network Entity that has a TCP/IP network address. It might be used by an iSCSI Node within that Network Entity for the connections within one of its iSCSI sessions. A Network Portal in a Target is identified by its IP address.

Network entities

Each IBMTSDS_StorageSystem instance is considered a separate network entity, associated to an instance of CIM_iSCSICapabilities through the CIM_ElementCapabilities association; it is an iSCSI network entity visible on an iSCSI network.
Instances of CIM_iSCSICapabilities contain the iSCSI specification versions and the authentication mechanisms that are supported by a network entity, which can be used to determine the capabilities of a storage array.

iSCSI nodes

iSCSI Node instances are represented by instances of CIM_SCSIProtocolController. This class can have many subclasses. The Name and NameFormat properties can be used to determine whether a CIM_SCSIProtocolController instance is an iSCSI Node instance.
All CIM_SCSIProtocolController instances that represent iSCSI Node instances are required to have a Name property whose value is an iSCSI Name value. Therefore, the NameFormat property must have a value of **3** (which maps to iSCSI Name).
On the IBM FlashSystem A9000 and A9000R arrays, iSCSI Node instances are represented by instances of IBMTSDS_iSCSIProtocolController. IBMTSDS_iSCSIProtocolController is a subclass of CIM_SCSIProtocolController, which is associated to the first-level system (IBMTSDS_StorageSystem) through CIM_SystemDevice.
For IBM FlashSystem A9000 and A9000R arrays, each IBMTSDS_StorageSystem entry has one Node entry for each IP Interface parameter. The iSCSI Name values for those nodes follow the IQN (iSCSI qualified name) format. The NameFormat and Name properties of the IBMTSDS_iSCSIProtocolController instances are shown bolded in the following example.

```
instance of IBMTSDS_iSCSIProtocolController { SystemCreationClassName =
"IBMTSDS_StorageSystem"; SystemName = "IBM.2810-6000095";
CreationClassName = "IBMTSDS_iSCSIProtocolController"; Name =
"iqn.2005-10.com.xivstorage:000095.M_7_1"; NameFormat = 3; ElementName =
"M_7_1"; DeviceID = "IBM.2810-6000095-100955"; Caption =
"iSCSI Node of IP Interface M_7_1"; HealthState = 0; };
```

iSCSI ports

Each iSCSI Node instance on an IBM FlashSystem A9000 and A9000R array is accessed through a single CIM_iSCSIProtocolEndpoint instance. The CIM_iSCSIProtocolEndpoint instances can be discovered by finding all iSCSI nodes, then following the CIM_SAPAvailableForElement association from the iSCSI node (IBMTSDS_iSCSIProtocolController) to CIM_iSCSIProtocolEndpoint.
To find all of the CIM_iSCSIProtocolEndpoint instances for a network entity, follow the CIM_HostedAccessPoint association from the IBMTSDS_StorageSystem instance to the CIM_iSCSIProtocolEndpoint instances.

Network Portal

There is a single network portal per IP Interface on an IBM FlashSystem A9000 and A9000R array, just as there is one iSCSI node and iSCSI port per IP Interface.
Network portals are represented by using the CIM_TCPProtocolEndpoint, CIM_IPProtocolEndpoint and CIM_EthernetPort classes, which are associated through the CIM_BindsTo association:

- The CIM_IPProtocolEndpoint instance contains the IP address of the network portal in the IPv4Address property. The IPv6Address property is not currently supported.
- The CIM_TCPProtocolEndpoint instance contains the TCP Port of the Network Portal in the PortNumber property.
- The CIM_EthernetPort instance contains the physical properties of the Network Portal in the PermanentAddress, DeviceID, FullDuplex, Speed, and OtherIdentifyingInfo properties.

iSCSI Target Ports methods

The following sections describe functional methods associated with the iSCSI Target Ports profile and object model, including uses of and parameters used by each method.

- **Creating an IP endpoint**
  You can use IBMTSDS_iSCSIConfigurationService.CreateIPProtocolEndpoint method to create an IP endpoint.
- **Modifying an IP endpoint**
  You can use IBMTSDS_iSCSIConfigurationService.ModifyIPProtocolEndpoint method to create an IP endpoint.
- **Deleting an IP endpoint**
  You can use the IBMTSDS_iSCSIConfigurationService. DeleteIPProtocolEndpoint method to delete an IP endpoint.

Creating an IP endpoint

You can use IBMTSDS_iSCSIConfigurationService.CreateIPProtocolEndpoint method to create an IP endpoint.

Procedure

1. Obtain the reference (CIMObjectPath) of an IBMTSDS_iSCSIConfigurationService instance. The instance is associated with the IBMTSDS_StorageSystem instance in which you receive the performance statistics by traversing the IBMTSDS_HostediSCSIConfigurationService association.
2. Invoke the IBMTSDS_iSCSIConfigurationService.CreateIPProtocolEndpoint method to create an IP endpoint with specified values of parameters ElementName, ModuleNumber, PortNumber, IPv4Address, IPv4Gateway, IPv4SubnetMask, and MTU.
3. After it is successfully completed, check the value of IPEndpoint representing the IP endpoint that was created.

- **Parameters**
  The following list describes the parameters of the IBMTSDS_iSCSIConfigurationService.CreateIPProtocolEndpoint method.

Parameters

The following list describes the parameters of the IBMTSDS_iSCSIConfigurationService.CreateIPProtocolEndpoint method.

ElementName *(required)*
: A meaningful name for the IP endpoint. It cannot be empty or null.

ModuleNumber *(required)*
: An identifier that indicates the module that contains the IP endpoint. It cannot be 0 or null.

PortNumber *(required)*
: Port identifier for this IP endpoint on this module. The value cannot be a value other than 1, 2, 3, or 4 because each module has four ports.

IPv4Address *(required)*
: IPv4 address for this port.

IPv4Gateway *(required)*
: IPv4 gateway for this port.

IPv4SubnetMask (required)
: IPv4 subnet mask for this port.

MTU
: MTU for this port. If it is null, default value 4500 is used.

IPEndpoint
: The reference to a point that was created.

Modifying an IP endpoint

You can use IBMTSDS_iSCSIConfigurationService.ModifyIPProtocolEndpoint method to create an IP endpoint.

Procedure

1. Obtain the reference (CIMObjectPath) of an IBMTSDS_iSCSIConfigurationService instance that is associated with the IBMTSDS_StorageSystem instance in which you receive the performance statistics by traversing the IBMTSDS_HostediSCSIConfigurationService association.
2. Obtain the reference (CIMObjectPath) of an IBMTSDS_IPEndPoint instance which is to be modified.
3. Invoke the IBMTSDS_iSCSIConfigurationService.ModifyIPProtocolEndpoint method to modify an IP endpoint with values of parameters ElementName, IPv4Address, IPv4Gateway, IPv4SubnetMask, MTU, and IPEndpoint.
4. After it is successfully completed, the IP endpoint will be modified.

- **Parameters**
  The following list describes the parameters of the IBMTSDS_iSCSIConfigurationService.ModifyIPProtocolEndpoint method.

Parameters

The following list describes the parameters of the IBMTSDS_iSCSIConfigurationService.ModifyIPProtocolEndpoint method.

> Note: The method ModifyIPProtocolEndpoint is not in the Common Information Model (CIM)/Storage Management Initiative Specification (SMI-S) schema but an IBM® extension.

ElementName
: A new meaningful name for the IP endpoint.

ModuleNumber
: An identifier that indicates the new module that contains the IP endpoint.

PortNumber
: The new port identifier for this IP endpoint. If it is set, the value cannot be a value other than 1, 2, 3, or 4 because each module has four ports.

IPv4Address
: A new IPv4 address for this port.

IPv4Gateway
: A new IPv4 gateway for this port.

IPv4SubnetMask
: A new IPv4 subnet mask for this port.

MTU
: A new MTU for this port.

IPEndpoint
: The reference to a point to modify.

Deleting an IP endpoint

You can use the IBMTSDS_iSCSIConfigurationService. DeleteIPProtocolEndpoint method to delete an IP endpoint.

Procedure

1. Obtain the reference (CIMObjectPath) of an IBMTSDS_iSCSIConfigurationService instance. Ensure that the reference is associated with the IBMTSDS_StorageSystem instance in which you receive the performance statistics by traversing the IBMTSDS_HostediSCSIConfigurationService association.
2. Obtain the reference (CIMObjectPath) of an IBMTSDS_IPEndPoint instance which is to be deleted.
3. Invoke the IBMTSDS_iSCSIConfigurationService.DeleteIPProtocolEndpoint method and set IPEndpoint as the reference obtained in step 2.
4. After it is successfully completed, the IP endpoint is deleted.

- **Parameters**
  The only input parameter that is required for the IBMTSDS_iSCSIConfigurationService. DeleteIPProtocolEndpoint method is a reference to the PortocolEndPoint instance that you want to delete. Only one PortocolEndPoint instance can be deleted at a time.

Parameters

The only input parameter that is required for the IBMTSDS_iSCSIConfigurationService. DeleteIPProtocolEndpoint method is a reference to the PortocolEndPoint instance that you want to delete. Only one PortocolEndPoint instance can be deleted at a time.

Note: The method DeleteIPProtocolEndpoint is not in the Common Information Model (CIM)/Storage Management Initiative Specification (SMI-S) schema but an IBM® extension.

Masking and Mapping profile

You can control which client hosts can see and access storage volumes by mapping and masking client hosts. The Storage Networking Industry Association (SNIA) Storage Management Initiative Specification (SMI-S) specifications define the Mapping and Masking profile to support this function.

LUN masking
: The process of configuring software in SAN nodes to determine which hosts have access to exported drives (volumes).
LUN masking for IBM FlashSystem® A9000 and A9000R systems is storage-based port mapping.

LUN mapping
: The process of creating a disk resource and defining its external access paths by using a LUN. The LUN is then mapped to an external ID descriptor. For example, a SCSI port or a target ID. To ensure uninterrupted data availability, map a logical volume to allow access from multiple ports, target IDs or both.

The IBM FlashSystem A9000 and A9000R system architecture for masking and mapping involves the following objects on two kinds of devices:

- Objects on IBM FlashSystem A9000 and A9000R systems:
  - Storage device I/O ports (FC ports and iSCSI ports)
  - Storage volumes
- Objects on the host system
  - Host ports (HBAs)
  - Hosts (host systems)
  - Clusters

Figure 1 provides an example of how IBM FlashSystem A9000 and A9000R systems use masking and mapping.

Figure 1. Masking and mapping physical model in IBM FlashSystem A9000 and A9000R systems

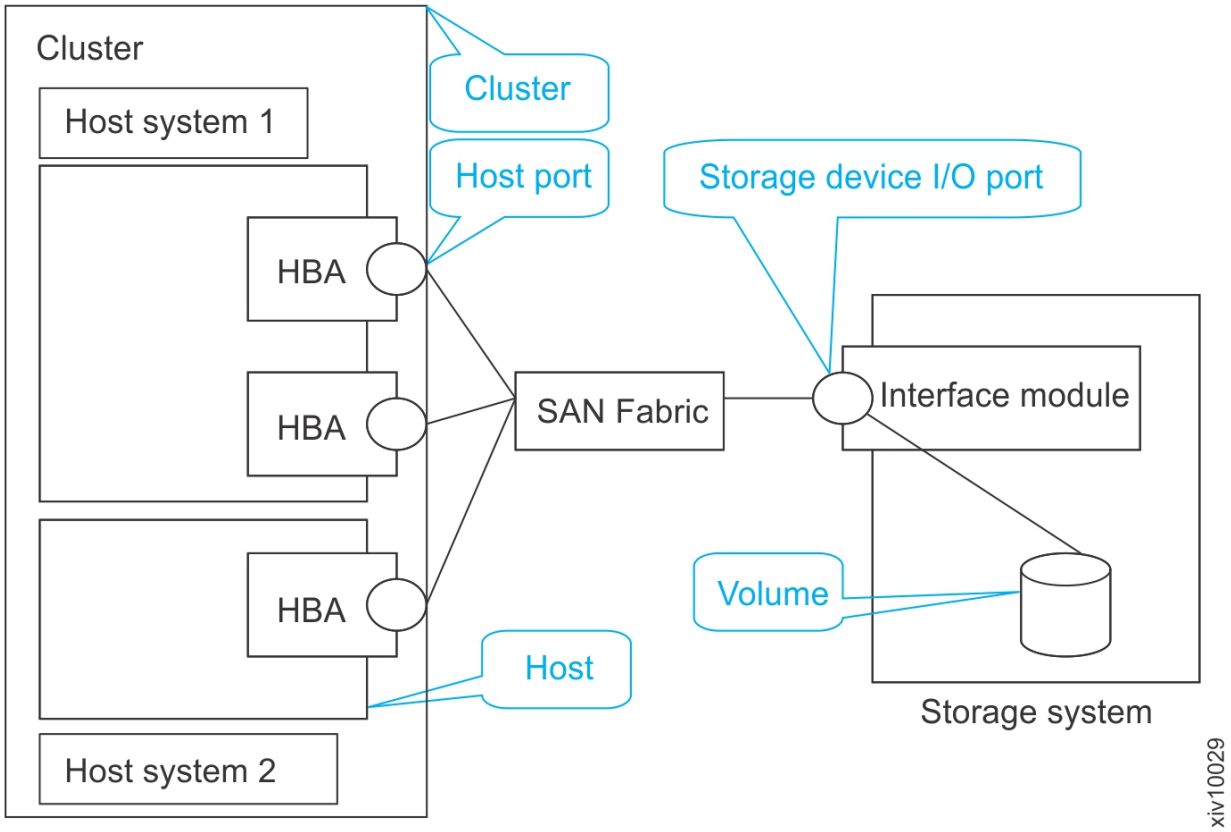

The following list defines each of the objects that are shown in Figure 1.

Storage device I/O ports
: Ports on IBM FlashSystem A9000 and A9000R systems that represent a Fibre Channel port or an iSCSI port that is used for host I/O operations. Storage device I/O ports are sometimes referred to as target ports.

Volumes
: Objects on IBM FlashSystem A9000 and A9000R systems that are mapped to host or a cluster. One volume cannot be mapped to multiple hosts that are not part of the same cluster.

Host ports
: Objects that represent a single Fibre Channel port or an iSCSI port on the host system. The architecture does not track host systems. If a host system has multiple host ports, a host port object must be created for each port.

Hosts
: Systems with a set of initiator ports that are used to access storage on the device. For each host system, you can create a host object, add host port objects, and then map the host to volumes through the IBM® Hyper-Scale Manager GUI, command-line interface (CLI), or Common Information Model (CIM) agent.

Clusters
: Groups of hosts. For IBM FlashSystem A9000 and A9000R systems, multiple hosts can see the same set of volumes. For a host system cluster, you can create a cluster object, add host objects to it, and then map the cluster object to multiple volumes through the IBM Hyper-Scale Manager GUI, CLI, or CIM agent. Clusters are required when more than one host can access the same volume.

After you set up the physical network and create a host object, add Fibre Channel ports or iSCSI ports to the host. Map the volumes to the host so that the volume can be used for I/O operations.

Mapping hosts to a volume

You can map a host to volumes. A host can contain several initiator ports that include Fibre Channel ports and iSCSI ports. In this case, create a host, add the initiator ports, and then map the host to a volume. Through this mapping, the host system is able to access the mapped volume through the included initiator ports.

Mapping clusters to a volume

You can map a cluster to a volume when several hosts must be mapped to the same volume. In this case, create a cluster and add hosts, and then map the cluster to volumes. Through this mapping, the volume is mapped to all hosts that are contained in the cluster. If a host is included in a cluster, it cannot be mapped to another volume separately.

See the *IBM FlashSystem A9000 and A9000R: Architecture, Implementation, and Usage* and the *IBM FlashSystem A9000 Product Overview* or *IBM FlashSystem A9000R Product Overview* for more information about the architecture of IBM FlashSystem A9000 and A9000R systems.

- **Masking and Mapping object model**
  The Masking and Mapping profile provides an interface to specify which hosts can see which volumes and through which target ports.
- **Masking and Mapping methods**
  The following sections describe functional methods associated with the Masking and Mapping profile and object model, including uses of and parameters used by each method.

Masking and Mapping object model

The Masking and Mapping profile provides an interface to specify which hosts can see which volumes and through which target ports.

For IBM FlashSystem A9000 and A9000R host and cluster mappings, specifying target ports for the view is not supported. Target ports cannot be selected when you configure host mappings, but can be configured by zoning configurations on the switch. The target ports that are involved in all views can be displayed by enumerating the instances of CIM_ProtocolControllerForPort class. See Figure 1.

Figure 1. Masking and mapping object model in SMI-S

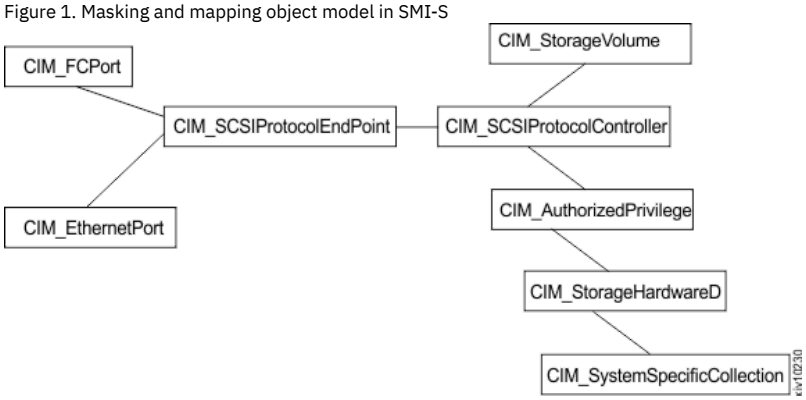

Table 1 describes the masking and mapping classes that are available. For more information about the masking and mapping object model, see the MaskingMapping.mof file. The MOF documentation is in the mof folder in the Common Information Model (CIM) agent installation directory.

Table 1. Masking and mapping classes

CIM class name	Description	IBM FlashSystem® A9000 and A9000R CIM class name
CIM_FCPort	A Fibre Channel port on the system	IBMTSDS_FCPort
CIM_EthernetPort	An Ethernet port on the system	IBMTSDS_EthernetPort
CIM_SCSIProtocolEndPoint	The Fibre Channel or iSCSI port that is used in the mapping	IBMTSDS_SCSIProtocolEndpoint IBMTSDS_iSCSIProtocolEndpoint
CIM_SCSIProtocolController	A logical entity that represents a host and the relationship to the volumes	IBMTSDS_SCSIProtocolController
CIM_StorageVolume	A storage volume on the system	IBMTSDS_SEVolume
CIM_AuthorizedPrivilege	A logical entity that represents the access permissions for a set of volumes	IBMTSDS_Privilege
CIM_StorageHardwareID	An initiator port on the host system	IBMTSDS_StorageHardwareID
CIM_SystemSpecificCollection	A logical entity which represents a collection of StorageHardwareID (host)	IBMTSDS_SystemSpecificCollection
CIM_Cluster	A logical entity which represents a collection of SystemSpecificCollection (hosts)	IBMTSDS_Cluster

The IBM FlashSystem A9000 and A9000R CIM agent provides a full implementation of masking and mapping subprofile including the following items:

- iSCSI port support
- StorageHardwareID instance creation and deletion
- HardwareIDCollection instance creation and modification
- ExposePaths and HidePaths methods

iSCSI port support

By default, IBM FlashSystem A9000 and A9000R ships equipped with six iSCSI ports. Three of the interface modules support iSCSI, with two ports in each module. You can refer to IBM FlashSystem A9000 and A9000R Redbook for detailed information. The IBM FlashSystem A9000 and A9000R CIM agent include the full iSCSI support by implementation of the iSCSI target ports profile and the masking and mapping profile.

StorageHardwareID instance manipulation

A StorageHardwareID instance represents the host-side initiator that logs in to a storage area network. The IBM FlashSystem A9000 and A9000R SMI-S provider supports host-side initiators that contain Fibre Channel or iSCSI host ports.

A StorageHardwareIDMangementService instance allows a CIM client to locate StorageHardwareID instances known to a storage array, and create, modify, and delete them.

- The CreateStorageHardwareID method creates a StorageHardwareID instance. On IBM FlashSystem A9000 and A9000R, a host is created which contains the specified initiator port.
- The DeleteStorageHardwareID method deletes a StorageHardwareID instance. On IBM FlashSystem A9000 and A9000R, the initiator port is removed from the host that contains it.

HardwareIDCollection instance manipulation

A HardwareIDCollection instance represents a collection of StorageHardwareID instances (host-side initiators) that log in to a storage area network.

A StorageHardwareIDMangementService instance allows a CIM client to locate HardwareIDCollection instances known to a storage array, create a HardwareIDCollection instance, and remove HardwareID instances from a HardwareIDCollection instance.

- The CreateHardwareIDCollection method creates a HardwareIDCollection instance with specified ElementName and Setting values.
- The AddHardwareIDsToCollection method adds StorageHardwareID instances to a HardwareIDCollection instance.

Note: The methods ModifyHardwareIDCollection, DeleteHardwareIDCollection, and RemoveHardwareIDsFromCollection are not in the CIM or SMI-S schema. These methods are an IBM® extension.

ExposePaths and HidePaths methods

An SCSIProtocolController instance represents a view of volumes that can be assigned access to a StorageHardwareID instance. You can use the ControllerConfigurationService method to locate SCSIProtocolController instances or the ExposePaths and HidePaths methods from StorageHardwareID instances to locate SCSIProtocolController instances.

A privilege determines the access rights between the StorageHardwareID instance and the SCSIProtocolController instance. The IBM FlashSystem A9000 and A9000R SMI-S provider supports read/write access. A privilege can be located by using the PrivilegeManagementService method.

Masking and Mapping methods

The following sections describe functional methods associated with the Masking and Mapping profile and object model, including uses of and parameters used by each method.

- **Creating a storage hardware ID**
  You can use the CreateStorageHardwareID method to create a StorageHardwareID instance. This method creates a host in the IBM FlashSystem A9000 or IBM FlashSystem A9000R device that contains the initiator port that is specified by StorageID value.
- **Deleting a storage hardware ID**
  Use the DeleteStorageHardwareID method to delete StorageHardwareID instances.
- **Creating a hardware ID collection**
  You can use the CreateHardwareIDCollection method to create HardwareIDCollection instances.
- **Adding hardware IDs to a collection**
  You can use the AddHardwareIDsToCollection method to add StorageHardwareID instances (representing initiator ports) into a SystemSpecificCollection instance (representing a host). The operation fails if the port already belongs to another host. The operation succeeds if the port already belongs to the specified host, but it has no effect.
- **Assigning volumes to a storage hardware ID**
  Use the IBMTSDS_ControllerConfigurationService. ExposePaths method to assign volumes to a StorageHardwareID instance representing a host or cluster in an IBM FlashSystem A9000 or IBM FlashSystem A9000R system.
- **Removing access to volumes from a storage hardware ID**
  The IBMTSDS_ControllerConfigurationService. HidePaths method is the inverse of the ExposePaths method; it is used to remove access to volumes from a StorageHardwareID instance.
- **Deleting an existing protocol controller**
  The IBMTSDS_ControllerConfigurationService. DeleteProtocolController method deletes an existing ProtocolController. In IBM FlashSystem A9000 and A9000R systems, the corresponding host and all the host ports that are contained by the host are deleted.
- **Example configuration procedures using the Masking and Mapping profile**
  The following sections provide example procedures to perform various configuration tasks using the Masking and Mapping profile, including the common tasks of configuring host and LUN mapping in an IBM FlashSystem A9000 or IBM FlashSystem A9000R system.

Creating a storage hardware ID

You can use the CreateStorageHardwareID method to create a StorageHardwareID instance. This method creates a host in the IBM FlashSystem A9000 or IBM FlashSystem A9000R device that contains the initiator port that is specified by StorageID value.

Parameters

The following list describes the parameters of the CreateStorageHardwareID method.

Setting
: An input parameter that represents the operating system of the port. It must be an instance of IBMTSDS_StorageClientSettingData. Each instance represents a different operating system. Only the values **`Standard`**, **`HPUX`**, and **`z/VM`** are supported by the IBM FlashSystem® A9000 and A9000R Common Information Model (CIM) agent. If it is left null, **`Standard`** is used by default.

IDType
: An input parameter that represents the type of port. To create a HardwareID instance representing a Fibre Channel port, specify 2; to create a HardwareID instance representing an iSCSI port, specify 5.

StorageID
: An input parameter that represents the ID of the port. If the IDType value is 2, specify it as the port's WWN; if the IDType value is 5, specify it as the port's IQN.

HardwareID
: An output parameter; the reference to the created StorageHardwareID instance.

Deleting a storage hardware ID

Use the DeleteStorageHardwareID method to delete StorageHardwareID instances.

Parameters

The following list describes the parameters of the DeleteStorageHardwareID method.

HardwareID
: An input parameter that represents a reference to the StorageHardwareID instance to delete.

Creating a hardware ID collection

You can use the CreateHardwareIDCollection method to create HardwareIDCollection instances.

Parameters

The following list describes the parameters of the CreateHardwareIDCollection method.

ElementName *(required)*
: An input parameter that represents a meaningful name for the HardwareIDCollection instance that is being created. The ElementName must be unique in the system. It cannot be a name that is already assigned to one of the other HardwareIDCollection instances in the system. If not specified, the Common Information Model (CIM) generates a random name.

HardwareIDs
: An input parameter that represents the ID of the ports that are contained by the HardwareIDCollection instance. If specified, each ID must be the WWN of a Fibre Channel port or the IQN of an iSCSI port. If not specified, an empty HardwareIDCollection instance is created.

HardwareIDCollection
: An output parameter that represents the reference to the created HardwareIDCollection instance.

Adding hardware IDs to a collection

You can use the AddHardwareIDsToCollection method to add StorageHardwareID instances (representing initiator ports) into a SystemSpecificCollection instance (representing a host). The operation fails if the port already belongs to another host. The operation succeeds if the port already belongs to the specified host, but it has no effect.

Parameters

The following list describes the parameters of the AddHardwareIDsToCollection method.

HardwareIDs *(required)*
: An input parameter that represents an array of strings that contain StorageID values of StorageHardwareID instances that become members of the HardwareIDCollection instance.

HardwareIDCollection *(required)*
: An input parameter that represents the reference to the SystemSpecificCollection instance into which the HardwareID values are added.

Assigning volumes to a storage hardware ID

Use the IBMTSDS_ControllerConfigurationService. ExposePaths method to assign volumes to a StorageHardwareID instance representing a host or cluster in an IBM FlashSystem A9000 or IBM FlashSystem A9000R system.

ExposePaths initiates the mapping and masking operation in one method call. It produces a list of volumes to a list of initiators, through one or more SCSIProtocolController instances (SPCs). It supports creating or modifying SPCs depending on different specified parameters.

The following table gives an overview of ExposePaths use cases and associated parameters and parameter values.

Table 1. ExposePaths use cases, parameters, and parameter values

Use cases	LUNames	Initiator PortIDs	Target PortIDs	Device Numbers	Device Accesses	ProtocolControllers (on input)
Create a view	Optional	Optional	NULL	Optional	Optional	NULL
Add LUNs to a view	Mandatory		NULL	Optional	Optional	contains a single SPC reference
Add initiator IDs to a view	NULL	Mandatory	NULL	NULL	NULL	contains a single SPC reference

Parameters (detailed)

The following list describes the parameters of the IBMTSDS_ControllerConfigurationService. ExposePaths method in more detail.

LUNames *(required)*
: A string array input parameter that represents the volumes to map to the SPC. They are not references to volume instances, but strings that match the Name property of the IBMTSDS_SEVolume (CIM_StorageVolume) instance, which is the WWN of an IBM FlashSystem A9000 or IBM FlashSystem A9000R volume.

InitiatorPortIDs *(required if StorageHardwareID has no volumes previously mapped)*
: A string array input parameter that represents the initiator ports to be added to an SPC. They are not references to StorageHardwareID instances, but the StorageID property of IBMTSDS_StorageHardwareID instance.

DeviceNumbers *(required)*
: A string array input parameter that represents the LUN IDs. Each item in this array must be a number 0 - 511. The number of items is the same as the number of items that are specified by the LUNames parameter. If not specified, the Common Information Model (CIM) assigns unused LUN IDs for the operation.

DeviceAccesses
: A unit16 array input parameter that represents the access rights to give to the StorageHardwareID instance and to the volumes specified in LUNames. Read/write access is supported. It must be an array the same size as that specified for LUNames, where each value is 2.

ProtocolControllers *(required if adding volumes to a StorageHardwareID that was already previously mapped to volumes)*
: As an input parameter, it represents the SCSIProtocolController instance to be modified. As an output parameter, it represents the SCSIProtocolController instance that is being modified or created.

Removing access to volumes from a storage hardware ID

The IBMTSDS_ControllerConfigurationService. HidePaths method is the inverse of the ExposePaths method; it is used to remove access to volumes from a StorageHardwareID instance.

The following table gives an overview of HidePaths use cases and associated parameters and parameter values.

Table 1. HidePaths use cases, parameters, and parameter values

Use case	LUNames	InitiatorPortIDs	TargetPortIDs	ProtocolControllers (on input)
Remove LUs from a view	Mandatory	NULL	NULL	contains a single SPC reference
Remove initiator IDs from a view	NULL	Mandatory	NULL	contains a single SPC reference
Hide full paths from a view	Mandatory	Mandatory	NULL	contains a single SPC reference

Parameters (detailed)

The following list describes the parameters of the IBMTSDS_ControllerConfigurationService. HidePaths method in more detail.

LUNames *(required for removing access to volumes)*
: A string array input parameter that represents the volumes to unmap from the StorageHardwareID instance. They are not references to volume instances, but are strings that match the Name property of the IBMTSDS_SEVolume (CIM_StorageVolume) instance.

InitiatorPortIDs *(required for removing access to initiator ports)*
: A string array input parameter that represents the initiator port that you want to remove from the SCSIProtocolController view.

ProtocolControllers *(required)*
: As an input parameter, it represents the SCSIProtocolController instance to be modified. As an output parameter, it represents the SCSIProtocolController instance that was modified or deleted.

Deleting an existing protocol controller

The IBMTSDS_ControllerConfigurationService. DeleteProtocolController method deletes an existing ProtocolController. In IBM FlashSystem A9000 and A9000R systems, the corresponding host and all the host ports that are contained by the host are deleted.

Parameters

The following list describes the parameters of the IBMTSDS_ControllerConfigurationService. DeleteProtocolController method.

ProtocolController *(required)*
: An input parameter that represents the ProtocolController instance to be deleted.

DeleteUnits
: An input parameter that represents the requirement to delete the volumes that are mapped to the ProtocolController instance. If true is specified, the mapped volumes are deleted.

Example configuration procedures using the Masking and Mapping profile

The following sections provide example procedures to perform various configuration tasks using the Masking and Mapping profile, including the common tasks of configuring host and LUN mapping in an IBM FlashSystem A9000 or IBM FlashSystem A9000R system.

Mapping FC and iSCSI ports to volumes (example):

1. Obtain the WWN of the host FC ports, the IQN of the host iSCSI ports, and the names of the volumes to be mapped to.

2. Use InvokeMethod on IBMTSDS_ControllerConfigurationService.ExposePaths, with InitiatorPortIDs set to the IDs of the host ports, LUNames set to the volume names, and DeviceAccesses set to 2.

Adding volumes to an existing SCSI protocol controller (SPC) (example):

1. Obtain the host reference and the names of the volumes to be added.
2. Find the IBMTSDS_SCSIProtocolController instance that matches the selected host.
3. Use InvokeMethod on IBMTSDS_ControllerConfigurationService.ExposePaths, with LUNames set to the volume names, DeviceAccesses set to 2, and ProtocolControllers set to the SCSIProtocolController instance name.

Adding initiator ports to an existing SPC (example):

1. Obtain the IDs of the initiator ports.
2. Find the IBMTSDS_SCSIProtocolController instance.
3. Use InvokeMethod on IBMTSDS_ControllerConfigurationService.ExposePaths, with InitiatorPortIDs set to the initiator port IDs, DeviceAccesses set to 2, and ProtocolControllers set to the SCSIProtocolController instance name.

Removing volumes from an SPC (example):

1. Obtain the SPC reference and names of the volumes to be removed.
2. Find the IBMTSDS_SCSIProtocolController instance that matches the selected host.
3. Use InvokeMethod on IBMTSDS_ControllerConfigurationService.HidePaths, with LUNames set to the volume names and ProtocolControllers set to the SCSIProtocolController instance name.

Removing initiator ports from an SPC (example):

1. Obtain the SPC reference and IDs of the initiator port.
2. Find the IBMTSDS_SCSIProtocolController instance.
3. Use InvokeMethod on IBMTSDS_ControllerConfigurationService.HidePaths, with InitiatorPortIDs set to the IDs of the initiator ports and ProtocolControllers set to the SCSIProtocolController instance name.

Indication profile

The Storage Management Initiative Specification (SMI-S) supports two types of indications: Lifecycle indications and alert indications.

Lifecycle indications are used to convey changes in the model and are concerned only with the creation, modification, or deletion of Common Information Model (CIM) instances. Alert indications are used to draw the attention of subscribing client applications to the occurrence of an event. Lifecycle indications are implemented in all IBM FlashSystem A9000 and A9000R CIM agent releases.

Supported lifecycle indications by class

Table 1 identifies the class of objects that can be monitored by using the specified indication type for IBM FlashSystem A9000 and A9000R arrays.

Table 1. Indication types and object classes

Class name	Supported lifecycle indications
IBMTSDS_SEVolume	IBMTSDS_InstCreation IBMTSDS_InstDeletion IBMTSDS_InstModification
IBMTSDS_VirtualPool	IBMTSDS_InstCreation IBMTSDS_InstDeletion IBMTSDS_InstModification
IBMTSDS_ProtocolControllerForSEUnit	IBMTSDS_InstCreation IBMTSDS_InstDeletion
IBMTSDS_SCSIProtocolController	IBMTSDS_InstCreation IBMTSDS_InstDeletion
IBMTSDS_SystemSpecificCollection	IBMTSDS_InstCreation IBMTSDS_InstDeletion
IBMTSDS_StorageHardwareID	IBMTSDS_InstCreation IBMTSDS_InstDeletion
IBMTSDS_iSCSIProtocolEndpoint	IBMTSDS_InstCreation IBMTSDS_InstDeletion

The lifecycle indications are sent following an IBM FlashSystem A9000 or IBM FlashSystem A9000R CIM agent configuration function being successfully completed. If the SMI-S client starts the CIM agent ExposePaths method to map a volume to a host, a related InstCreation indication with the new

IBMTSDS_ProtocolControllerForSEUnit instance is sent from the provider.

In addition, the CIM agent periodically fetches events and sends related lifecycle indications. If a StorageVolume instance is created on an array through the IBM FlashSystem A9000 or IBM FlashSystem A9000R GUI, an SMI-S client receives a related InstCreation indication with the new IBMTSDS_SEVolume instance from the provider. By default, the IBM FlashSystem A9000 and A9000R event fetch interval is 30 seconds.

Alert indications

The IBM FlashSystem A9000 and A9000R CIM agent periodically fetches non-internal events and sends one alert indication for each event. The alert indication comes in the form of class IBMTSDS_AlertIndication. By default, the IBM FlashSystem A9000 and A9000R events fetching interval is 30 seconds.

Note: The IBM FlashSystem A9000 and A9000R CIM agent has a limitation of values of properties:

- IndicationTime of IBMTSDS_InstCreation class
- IndicationTime of IBMTSDS_InstDeletion class
- IndicationTime of IBMTSDS_InstModification class
- EventTime of IBMTSDS_AlertIndication class
- IndicationTime of IBMTSDS_AlertIndication class

The above properties are representing a timestamp in format `yyyymmddhhmmss.mmmmmmsutc`, where:

- `yyyy` is the four-digit year
- `mm` is the month within the year (starting with 01)
- `dd` is the day within the month (starting with 01)
- `hh` is the hour within the day (24-hour clock, starting with 00)
- `mm` is the minute within the hour (starting with 00)
- `ss` is the second within the minute (starting with 00)
- `mmmmmm` is the microsecond within the second (starting with 000000)
- `s` is a "+" or "-", indicating that the value is a timestamp, and indicating the sign of the UTC (Universal Coordinated Time) correction field. A "+" is used for time zones east of Greenwich, and a "-" is used for time zones west of Greenwich.
- `utc` is the offset from UTC in minutes (using the sign indicated by s)

Fields other than `utc` are correct which are representing the UTC timestamp and the `utc` field as 000. Occasionally the `utc` property has values other than 000 (such as 540). Values other than 000 are be ignored.

Replication Services profile

The Replication Services profile allows a storage system to copy data from a source element to a target element. The copy operations can be initiated on elements from the same storage system or across a connection to a different storage system.

Elements that are used in a copy operation can be grouped to facilitate the copy operation on many elements at the same time. Furthermore, the elements of a group can be declared as consistent.

Two types of synchronization views are supported. A target element can be synchronized to the current view of the source element, or it can be synchronized to a point-in-time view. Snapshots and clones represent a point-in-time view, while a mirror represents a current view.

Two copy modes are supported: synchronous and asynchronous. In synchronous mode, the writer waits for acknowledgment that a write to the source element was processed by the target element. Until it receives this acknowledgment, it does not accept another I/O from the host.

In asynchronous mode, the writer does not wait for this acknowledgment and continues processing I/Os, enabling writes to be sent to the target element later.

The Replication Services profile supports local and remote replication. Local replication specifies that both the source elements and the target elements are contained in a single managed system, such as an array platform.

Remote replication specifies that the source elements and the target elements are contained in separate systems.

For remote replication, the client may interact with both the source system and the target system, but the client only invokes the replication methods to a single replication service.

- **Replication Services object model**
  Data can be copied from a source element or group of elements to a target element. The data can be copied within the same storage system, or across a connection to a different storage system.
- **Replication Services methods**
  The following sections describe functional methods associated with the Replication Services profile and object model, including uses of and parameters used by each method.
- **Replication Services indications**
  The IBM FlashSystem A9000 and A9000R Common Information Model (CIM) agent supports the following Replication Services profile indications.

Replication Services object model

Data can be copied from a source element or group of elements to a target element. The data can be copied within the same storage system, or across a connection to a different storage system.

See Figure 1 and Figure 2.

Figure 1. Replication Services (Local) SMI-S model for IBM FlashSystem A9000 and A9000R systems

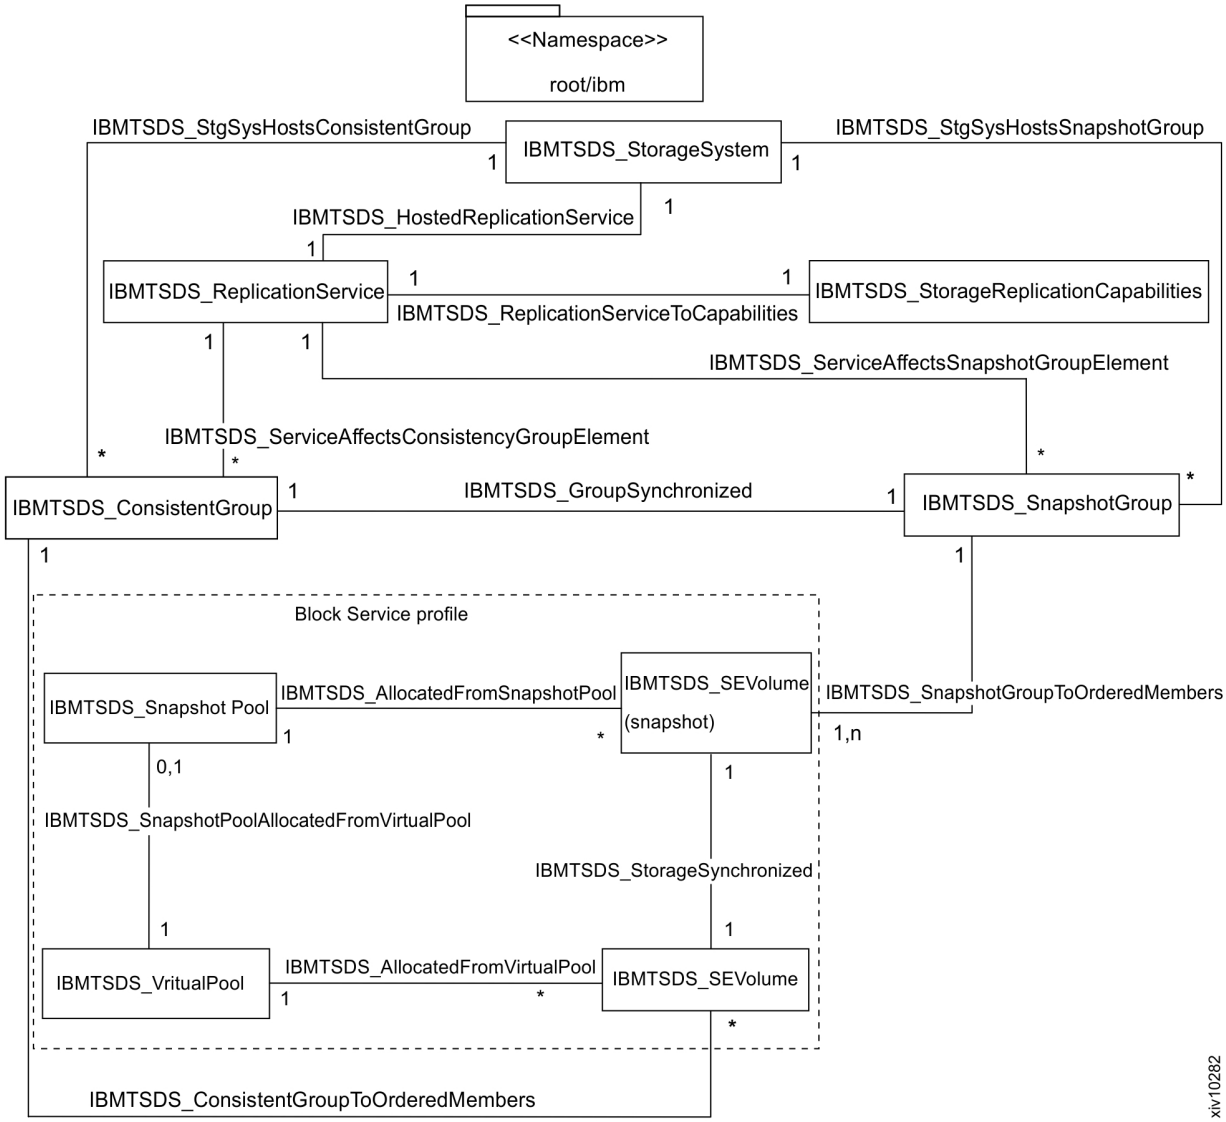

Figure 2. Replication Services (Remote) SMI-S model for IBM FlashSystem A9000 and A9000R systems

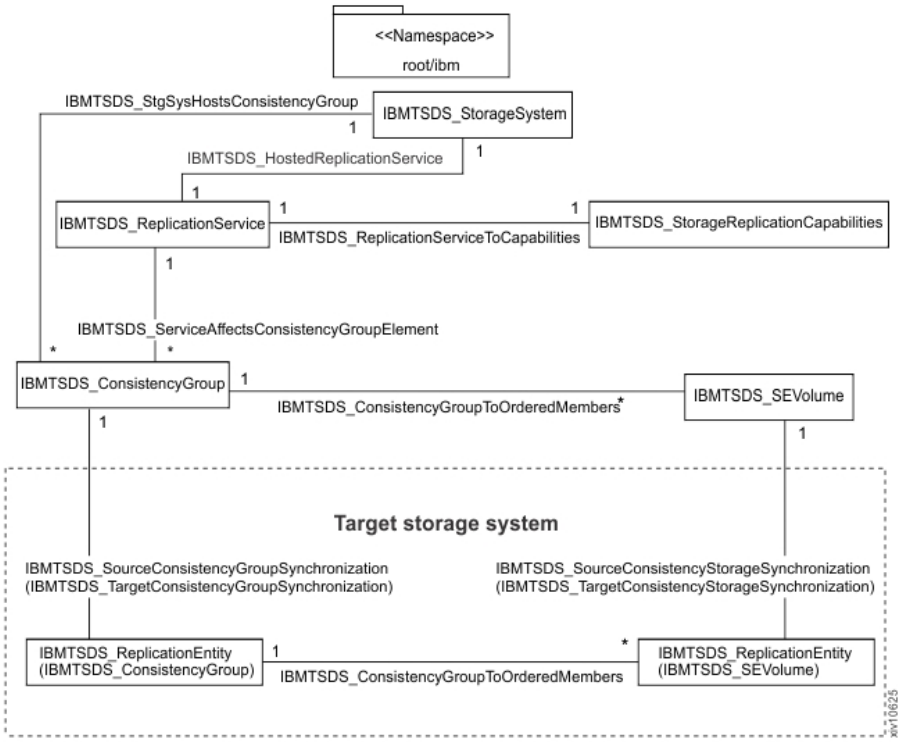

Replication Service capabilities

Use IBM FlashSystem A9000 and A9000R terminology to map to Storage Management Initiative Specification (SMI-S) specific terminology:

Table 1. Mapping IBM FlashSystem A9000 and A9000R terminology to SMI terminology

IBM FlashSystem A9000 and A9000R Terminology	SMI Terminology	Description	IBM FlashSystem A9000 and A9000R CIM Agent support
Snapshot	Asynchronous Snapshot Local	A "Point-in-Time", associated virtual copy of the source element. The target element enables visibility into a session.	version 12.0 or later
Volume copy	Synchronous Clone Local	A full, "Point-In-Time", unassociated local copy of the source element.	version 12.0 or later
Synchronous Remote Mirroring	Synchronous Mirror Remote	A synchronized remote copy of the source element.	version 12.0.1 or later
Asynchronous Remote Mirroring	Asynchronous Mirror Remote	An asynchronous remote copy of the source element.	version 12.0.2 or later

The single instance of the class ReplicationServiceCapabilities and its methods describe the various capabilities of the service. You can examine the ReplicationServiceCapabilities instance and invoke its methods to determine the specific capabilities of a replication service implementation.

Table 2. Replication Service methods

Method	Description
ConvertSyncTypeToReplicationType	Translates CopyType, Mode, and Local/Remote to the corresponding ReplicationType
ConvertReplicationTypeToSyncType	Translates ReplicationType to the corresponding CopyType, Mode, and Local/Remote
GetSupportedFeatures	Determines supported features
GetSupportedGroupFeatures	Determines supported group features
GetSupportedCopyStates	Determines the supported copy states
GetSupportedGroupCopyStates	Determines supported group copy states
GetSupportedWaitForCopyStates	Determines supported wait for copy states
GetSupportedConsistency	Determines supported consistency
GetSupportedOperations	Determines supported operations
GetSupportedGroupOperations	Determines supported group operations
GetSupportedListOperations	Determines supported list operations
GetSupportedMaximum	Determines supported maximum
GetDefaultConsistency	Determines default consistency
GetDefaultGroupPersistency	Determines default group persistency
GetSynchronizationSupported	For the supplied element, this method returns the supported synchronization operations in a series of parallel output arrays.
GetSupportedOperationsForSynchronization	For the supplied synchronized association, this method returns the supported operations in a series of parallel output arrays.

Replication Services methods

The following sections describe functional methods associated with the Replication Services profile and object model, including uses of and parameters used by each method.

- **Group management**
  The CIM agent supports group manipulation of consistency groups and snapshot groups.
- **Creating a consistency group**
  The IBMTSDS_ReplicationService. CreateGroup method creates a consistency group.
- **Deleting a consistency group**
  The IBMTSDS_ReplicationService. DeleteGroup method deletes an existing consistency group.
- **Adding volumes to a consistency group**
  The IBMTSDS_ReplicationService. AddMembers method adds volumes to a consistency group.
- **Removing volumes from a consistency group**
  The IBMTSDS_ReplicationService. RemoveMembers method removes volumes from a consistency group.
- **Replication management**
  IBM FlashSystem A9000 and A9000R support local replication manipulation, including local snapshot and clone, and also supports remote replication manipulation, including volume and consistency group mirroring.
- **Creating a snapshot or clone**
  The IBMTSDS_ReplicationService. CreateElementReplica method creates a snapshot of a volume, duplicate snapshot of a snapshot, clone of a volume, and clone of a snapshot.
- **Creating a snapshot group for a consistency group**
  The IBMTSDS_ReplicationService. CreateGroupReplica method creates a snapshot group for a consistency group.
- **Restoring a snapshot to a volume or manipulating a mirror consistency group**
  For local replication service, use the IBMTSDS_ReplicationService. ModifyReplicaSynchronization method to restore a snapshot to a volume, restore a snapshot group to a consistency group, delete a snapshot, and delete a snapshot group. For remote replication service, use the IBMTSDS_ReplicationService. ModifyReplicaSynchronization method to manipulate the mirror consistency group (for example, to split the mirrored consistency group or set up failover for the mirrored consistency group).
- **Modifying a list of synchronizations using a batch operation**
  The IBMTSDS_ReplicationService. ModifyListSynchronization method is a batch operation of the IBMTSDS_ReplicationService. ModifyReplicaSynchronization method, which modifies a list of synchronizations.
- **Retrieving target elements**
  The IBMTSDS_ReplicationService. GetAvailableTargetElements method retrieves all of the candidate target elements for the supplied source element.
- **Retrieving snapshot or mirror relationships**
  The IBMTSDS_ReplicationService. GetReplicationRelationships method retrieves all of the snapshot or mirror relationships on an IBM FlashSystem A9000 or IBM FlashSystem A9000R device.
- **Retrieving strings of references to snapshots, snapshot groups, or mirror relationships**
  Use the IBMTSDS_ReplicationService. GetReplicationRelationshipInstances to retrieve strings of references to snapshots, snapshot groups, or mirror relationships.
- **Retrieving peer systems**
  Use the IBMTSDS_ReplicationService. GetPeerSystems method to retrieve (or start a job to retreive) all of the peer systems. A peer system is a system that is known and visible to the replication service. This method will return all the connected IBM FlashSystem A9000 or IBM FlashSystem A9000R systems.
- **Creating new storage objects that are replicas of specified source storage objects**
  Use the IBMTSDS_ReplicationService. CreateGroupReplicaFromElements method to create (or start a job to create) new storage objects that are replicas of the specified source storage objects (SourceElements).

Group management

The CIM agent supports group manipulation of consistency groups and snapshot groups.

See Table 1, Figure 1, and Table 2.

Table 1. Group management classes

IBM FlashSystem A9000 and A9000R	Storage Management Initiative Specification (SMI-S) class	IBM FlashSystem A9000 and A9000R CIM class
Consistency group	CIM_ReplicationGroup	IBMTSDS_ConsistencyGroup
Snapshot group	CIM_ReplicationGroup	IBMTSDS_SnapshotGroup
Association between Consistency group and volumes in it	CIM_OrderedMemberOfCollection	IBMTSDS_ConsistencyGroupTo OrderedMembers
Association between Snapshot group and snapshots in it	CIM_OrderedMemberOfCollection	IBMTSDS_SnapshotGroupTo OrderedMembers

Figure 1. Sample local group information retrieval

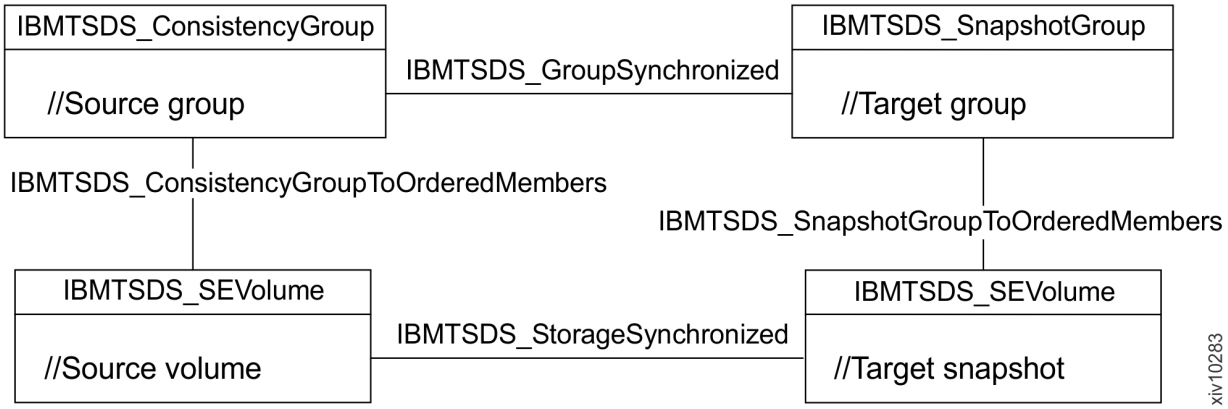

Table 2. Extrinsic methods for group management

Method	Description
CreateGroup	Creates a consistency group
DeleteGroup	Deletes an existing consistency group
AddMembers	Adds volumes to a consistency group
RemoveMembers	Removes volumes from a consistency group

Creating a consistency group

The IBMTSDS_ReplicationService. CreateGroup method creates a consistency group.

Parameters

The following list describes the parameters of the IBMTSDS_ReplicationService. CreateGroup method.

GroupName
: As an input, refers to the name of the group that is created. If not specified, a random name is generated.

Members
: As an input, refers the members to be added to the group that is created. If not specified, an empty group is created.

ReplicationGroup
: As an output, refers to the group that is created.

Deleting a consistency group

The IBMTSDS_ReplicationService. DeleteGroup method deletes an existing consistency group.

Parameters

The following list describes the parameters of the IBMTSDS_ReplicationService. DeleteGroup method.

RemoveElements
: As an input, a value of true indicates that members in this group are removed before you delete the group. If one or more elements in the group are in a replication relationship, its value is ignored.

ReplicationGroup
: As an input, refers to the consistency group that you want to delete.

Adding volumes to a consistency group

The IBMTSDS_ReplicationService. AddMembers method adds volumes to a consistency group.

Parameters

The following list describes the parameters of the IBMTSDS_ReplicationService. AddMembers method.

ReplicationGroup
: As an input, refers to the consistency group to which members are added.

Members
: As an input, refers to volumes to be added to the group. All specified volumes must be in the same pool of the consistency group.

Removing volumes from a consistency group

The IBMTSDS_ReplicationService. RemoveMembers method removes volumes from a consistency group.

Parameters

The following list describes the parameters of the IBMTSDS_ReplicationService. RemoveMembers method.

ReplicationGroup
: As an input, refers to the consistency group from which members are removed.

Members
: As an input, refers to volumes to be removed from the group.

DeleteOnEmptyElement
: As an input, a value of true indicates that the consistency group is deleted if all members are removed.

Replication management

IBM FlashSystem A9000 and A9000R support local replication manipulation, including local snapshot and clone, and also supports remote replication manipulation, including volume and consistency group mirroring.

See the following tables and figures for information about replication management:

Table 1. Replication management classes

IBM FlashSystem A9000 and A9000R	SMI class	IBM FlashSystem A9000 and A9000R CIM class
Volume, snapshot	CIM_StorageVolume	IBMTSDS_SEVolume
Remote volume / consistency group	CIM_ReplicationEntity	IBMTSDS_ReplicationEntity
Association between volume and its snapshot	CIM_StorageSynchronized	IBMTSDS_StorageSynchronized
Association between consistency group and its snapshot group	CIM_GroupSynchronized	IBMTSDS_GroupSynchronized
Association between mirrored volumes	CIM_StorageSynchronized	IBMTSDS_SourceConsistencyStorageSynchronized/IBMTSDS_TargetConsistencyStorageSynchronized
Association between mirrored consistency groups	CIM_GroupSynchronized	IBMTSDS_SourceConsistencyGroupSynchronized/IBMTSDS_TargetConsistencyGroupSynchronized

See Figure 1.

Figure 1. Sample remote group information retrieval

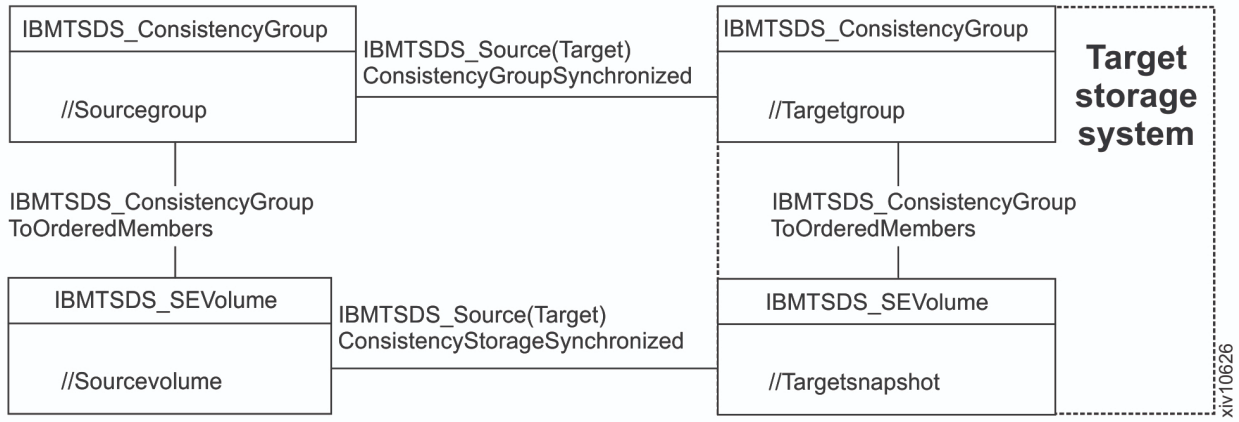

Note: The IBMTSDS_SourceConsistencyStorageSynchronized and IBMTSDS_TargetConsistencyStorageSynchronized classes are association classes which associate the source volume with the target volume.
They are essentially the same association class, but differ in how they get data from the Primary IBM FlashSystem A9000 or IBM FlashSystem A9000R and the Secondary IBM FlashSystem A9000 or IBM FlashSystem A9000R respectively; see Figure 2.

At the same time, the IBMTSDS_SourceConsistencyGroupSynchronized and IBMTSDS_TargetConsistencyGroupSynchronized classes associate the source consistency group with the target consistency group.

Figure 2. Association classes for mirrored volumes and consistency groups

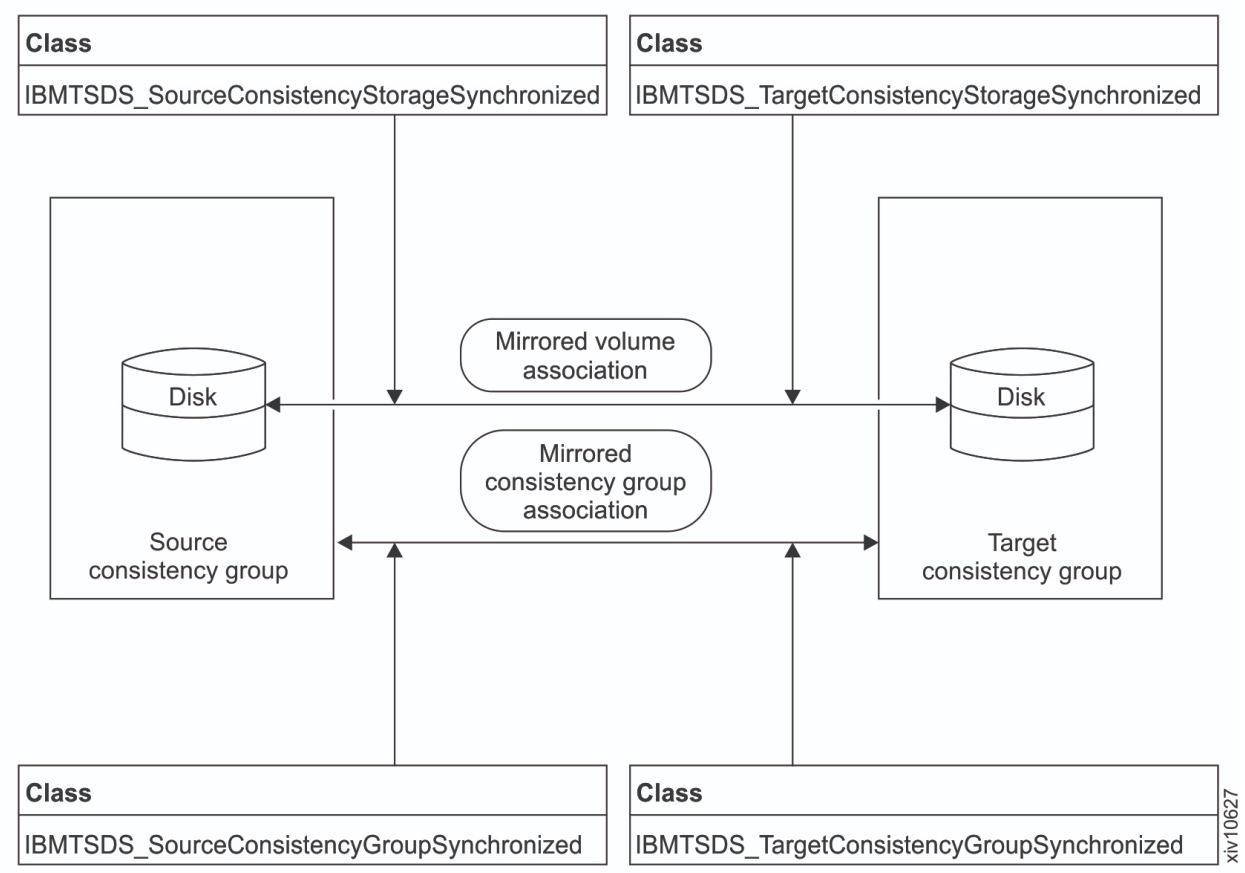

Table 2. Extrinsic methods for replication management

Method	Description
CreateElementReplica	Create a snapshot of a volume, or a duplication snapshot of a snapshot, or a clone of a volume.
CreateGroupReplica	Create a snapshot group of a consistency group.
ModifyReplicaSynchronization	For local replication service, restore a snapshot or a snapshot group; or delete a snapshot or a snapshot group.
ModifyListSynchronization	Restore a list of snapshots or snapshot groups; or delete snapshots or snapshot groups.
GetAvailableTargetElements	Return snapshots of a specified volume or snapshot.
GetPeerSystems	Get all of the peer systems.
CreateGroupReplicaFromElements	Create consistency group and volumes mirror.
GetReplicationRelationships	Get all of available synchronization relationships.
GetReplicationRelationshipInstances	Get all of available synchronization relationship instances as strings of references.

Creating a snapshot or clone

The IBMTSDS_ReplicationService. CreateElementReplica method creates a snapshot of a volume, duplicate snapshot of a snapshot, clone of a volume, and clone of a snapshot.

Parameters

The following list describes the parameters of the IBMTSDS_ReplicationService. CreateElementReplica method.

ElementName
: As an input, refers the name for the volume or clone being created. If not specified, a random name is generated.

SyncType
: As an input, refers the type of replica to be created. To create a snapshot, specify 7; to create a clone, specify 8.

Mode
: As an input, refers the mode of replica to be created. Only 3 is supported, which means an asynchronous replica is created if specified.

SourceElement
: As an input, refers the source volume or snapshot for which the replica is created.

TargetElement
: As an input, refers to a target element to use if specified. This parameter cannot be specified when you are creating a snapshot. As an output, refers to the replica that is created.

TargetPool
: As an input, refers the pool in which the replica is being created. For snapshot creation, if specified, it must be an instance of IBMTSDS_SnapshotPool which is in the same pool of the source element. For clone creation, if specified, it can be an instance of IBMTSDS_VirtualPool representing any pool in the device.

If not specified, the Common Information Model (CIM) selects the first available pool on the device to create the clone.

WaitForCopyState
As an input, refers to the copy state the replica must reach before the method returns. Only 4 is supported if specified.

Creating a snapshot group for a consistency group

The IBMTSDS_ReplicationService. CreateGroupReplica method creates a snapshot group for a consistency group.

Parameters

The following list describes the parameters of the IBMTSDS_ReplicationService. CreateGroupReplica method.

RelationshipName
As an input, refers to the name of the snapshot group that is being created. If not specified, a random name is generated.
SyncType
As an input, refers the type of replica to be created. Only 7 is supported if specified because this method can create a snapshot of a consistency group.
Mode
As an input, refers the mode of replica to be created. Only 3 is supported if specified, which means an asynchronous snapshot group is created.
SourceGroup
As an input, refers to a consistency group for which the snapshot group is being created.
TargetGroup
As an output, refers to the snapshot group that is being created.
Consistency
As an input, refers to the group consistency. Only 3 is supported if specified.
Synchronization
As an output, refers to the created association between the source consistency group and the snapshot group that is created.
TargetPool
As an input, refers to the pool in which the replica is being created. The pool must be the same as the pool of the SourceGroup if specified.
WaitForCopyState
As an input, refers to the copy state the replica must reach before the method returns. Only 4 is supported if specified.

Restoring a snapshot to a volume or manipulating a mirror consistency group

For local replication service, use the IBMTSDS_ReplicationService. ModifyReplicaSynchronization method to restore a snapshot to a volume, restore a snapshot group to a consistency group, delete a snapshot, and delete a snapshot group. For remote replication service, use the IBMTSDS_ReplicationService. ModifyReplicaSynchronization method to manipulate the mirror consistency group (for example, to split the mirrored consistency group or set up failover for the mirrored consistency group).

Parameters

The following list describes the parameters of the IBMTSDS_ReplicationService. ModifyReplicaSynchronization method.

Operation
As an input, refers to the type of modification to be made to the replica. The following values/modifications are supported: 4 (Activate), 5 (AddSyncPair), 7 (Deactivate), 8 (Detach), 10 (Failover), 12 (Fracture), 13 (RemoveSyncPair), 14 (ReSync Replica), 15 (Restore from Replica), 19 (Return To ResourcePool) and 20 (Reverse Roles), and 21 (Split).
Synchronization
As an input, refers to the replica synchronization to be modified. For local, it must be an instance of IBMTSDS_GroupSynchronized or IBMTSDS_StorageSynchronized class. For remote, it must be an instance of IBMTSDS_SourceConsistencyGroupSynchronized or IBMTSDS_SourceConsistencyStorageSynchronized.
WaitForCopyState
As an input, refers to the copy state the replica must reach before the method returns. The values 3, 4, 8, and 10 are supported, if specified. Different operations correspond with different WaitForCopyState values, as shown in Table 1.

Table 1. Operations, operation descriptions, and corresponding WaitForCopyState states

Operation	Operation Description (from profile)	Corresponding WaitForCopyState Value
Activate	Activate an Inactive or Prepared StorageSynchronized association.	Unsynchronized
Detach	Remove the association between the source and target elements. Does not delete the target element.	N/A
Failover	Enable the read and write operations from the host to the target element. This operation is useful for situations when the source element is unavailable.	Failedover
Fracture/Split	Separate the target element from the source element.	Inactive
Resync Replica	Resynchronize a fractured target element. Or, from a Broken or Aborted relationship. To continue from the Broken state, the problem should be corrected first before resyncing the replica. Also, to continue from the Aborted state.	Synchronized
Reverse Roles	Switch the source and the target element roles. The source element may need to be Read Only. See GetSupportedFeatures in capabilities.	Unsynchronized
Add/Remove SyncPair	Add/Remove mirrored volumes into/from mirrored consistency group.	Synchronized

Note: The ReverseRoles and ResyncReplica operations may take some time to execute, so a Job will be returned (4096).

Modifying a list of synchronizations using a batch operation

The IBMTSDS_ReplicationService. ModifyListSynchronization method is a batch operation of the IBMTSDS_ReplicationService. ModifyReplicaSynchronization method, which modifies a list of synchronizations.

For a list of applicable parameters, see Restoring a snapshot to a volume or manipulating a mirror consistency group.

Retrieving target elements

The IBMTSDS_ReplicationService. GetAvailableTargetElements method retrieves all of the candidate target elements for the supplied source element.

Parameters

The following list describes the parameters of the IBMTSDS_ReplicationService. GetAvailableTargetElements method.

SourceElement
: As an input, refers to the source volume. It must be an instance of IBMTSDS_SEVolume class.

SyncType
: As an input, refers to type of target elements. Only 7 is supported if specified.

Mode
: As an input, refers to the mode of target elements. Only 3 is supported if specified.

Candidates
: As an output, refers to the target elements of the source volume.

Retrieving snapshot or mirror relationships

The IBMTSDS_ReplicationService. GetReplicationRelationships method retrieves all of the snapshot or mirror relationships on an IBM FlashSystem A9000 or IBM FlashSystem A9000R device.

Parameters

The following list describes the parameters of the IBMTSDS_ReplicationService. GetReplicationRelationships method.

Type
: As an input, refers to the replication type. To get snapshots and mirrored volumes, specify 2; to get snapshot groups, specify 3; if not specified, all snapshots and snapshot groups are returned.

SyncType
: As an input, refers to type of target elements. To get snapshots and snapshot groups, specify 7. To get mirror relationships, specify 7; if no value is specified, all snapshots, snapshot groups, and mirror relationships are returned.

Mode
: As an input, refers to mode of target elements. 2 and 3 are supported if specified.

Locality
: As an input, refers to locality of target elements. To get mirrored volumes and consistency groups, specify 3. To get snapshots and snapshot groups, specify 2. If no value is specified, all snapshots, snapshot groups, and mirror relationships are returned.

CopyState
: As an input, refers to copy state of target elements. Only 4 is supported if specified.

Synchronizations
: As an input, refers to references of the returned snapshots, snapshot groups, or mirror relationships.

Retrieving strings of references to snapshots, snapshot groups, or mirror relationships

Use the IBMTSDS_ReplicationService. GetReplicationRelationshipInstances to retrieve strings of references to snapshots, snapshot groups, or mirror relationships.

Retrieving peer systems

Use the IBMTSDS_ReplicationService. GetPeerSystems method to retrieve (or start a job to retreive) all of the peer systems. A peer system is a system that is known and visible to the replication service. This method will return all the connected IBM FlashSystem A9000 or IBM FlashSystem A9000R systems.

Creating new storage objects that are replicas of specified source storage objects

Use the IBMTSDS_ReplicationService. CreateGroupReplicaFromElements method to create (or start a job to create) new storage objects that are replicas of the specified source storage objects (SourceElements).

If 0 is returned, the function completed successfully, and no ConcreteJob instance is created. If 4096/0x1000 is returned, a ConcreteJob is started, and a reference to it is returned in the Job output parameter. This method combines the functionality of the CreateGroup and CreateGroupReplica methods, in one call. This method creates mirrors for volumes and consistency groups and then adds the mirrored volumes into the mirrored consistency groups.

Parameters

The following list describes the parameters of the IBMTSDS_ReplicationService. CreateGroupReplicaFromElements method.

SyncType
: As an input, refers to the type of target elements. Only 6 is supported.

Mode
: As an input, refers to the mode of target elements. 2 and 3 are supported, if specified.

SourceGroup
: As an input, refers to a source group. It must be an instance of IBMTSDS_ConsistencyGroup class.

SourceElements
: As an input, refers to the source volumes. It must be an instance of IBMTSDS_SEVolume class.

SourceGroupName
: As an input, refers to the name of the source group to be created. As an output, refers to the name of the source group that is being created.

TargetGroup
: As an input, refers to a target group. It must be an instance of IBMTSDS_ConsistencyGroup class. If not specified, a target group will be created in the target pool, which is specified by the TargetPool parameter. TargetGroup and TargetPool can't be NULL at the same time.

TargetGroupName
: As an input, refers to the name of the target group to be created.

Consistency
: As an input, only 3 is supported, if specified.

TargetPool
: As an input, refers to the target pool that puts the target elements (the replicas). It must be an instance of IBMTSDS_VirtualPool. If not specified, it will be gotten from the TargetGroup parameter. If TargetGroup and TargetPool are specified at the same time, and TargetGroup is not in the TargetPool parameter, TargetGroup will be moved to it.

WaitForCopyState
: As an input, refers to the copy state the replica must reach before the method returns. Only 3 is supported, if specified.

Job
: As an output, refers to the job.

Replication Services indications

The IBM FlashSystem A9000 and A9000R Common Information Model (CIM) agent supports the following Replication Services profile indications.

Table 1. Replication Services profile indications

Indication	Description
IBMTSDS_StorageSynchronized InstCreation	After a snapshot is created or snapshot group is created
IBMTSDS_GroupSynchronized InstCreation	After a snapshot group is created
IBMTSDS_GroupSynchronized InstDeletion	After a snapshot group is deleted
IBMTSDS_StorageSynchronized InstDeletion	After a snapshot or snapshot group is deleted
IBMTSDS_ConsistencyGroupInstDeletion	After a consistency group is deleted
IBMTSDS_ConsistencyGroupInstCreation	After a consistency group is created
IBMTSDS_ConsistencyGroupInstModification	After a consistency group is renamed
IBMTSDS_ConsistencyGroupToOrderedMembersInstCreation	After volume is added/removed into/from consistency group
IBMTSDS_ConsistencyGroupToOrderedMembersInstDeletion	After all volumes are removed from the consistency group
IBMTSDS_SourceConsistencyGroupSynchronizedInstDeletion	After the mirror relationship is deleted
IBMTSDS_SourceConsistencyGroupSynchronizedInstCreation	After the mirror relationship is created
IBMTSDS_TargetConsistencyGroupSynchronizedInstCreation	After the mirror relationship is created
IBMTSDS_SourceConsistencyGroupSynchronizedInstModification	After the mirror relationship designation is changed
IBMTSDS_TargetConsistencyGroupSynchronizedInstModification	After the mirror relationship designation is changed

Job Control profile

In some profiles, like the replication services profile, some of the methods described may take some time to execute. In these cases, a mechanism is needed to handle asynchronous execution of the method as a 'job.' The Job Control profile defines the constructs and behavior for job control.

When the Job Control profile is implemented, and a client executes a method that executes asynchronously, a reference to an instance of ConcreteJob is returned and the return value for the method is set to **`Method parameters checked - job started (4096)`**. The ConcreteJob instance allows the progress of the method to be checked, and instance indications can be used to subscribe for job completion.

Although there are many methods, such as format volume and copy volume that are asynchronous jobs, in this work item, we only implement the job control for the volume/consistency group mirror.

In the IBM FlashSystem A9000 and A9000R Common Information Model (CIM) provider, the methods CreateGroupReplicaFromElements and ModifyReplicaSynchronization will output a job that will be associated to the elements whose references are created or modified as a side-effect of the job's execution via the AffectJobElement association.

The lifetime of a completed job instance, and thus the AffectedJobElement association to the appropriate Element instance, is currently implementation dependent. But, the set of AffectedJobElement associations to the Input and Output elements present when the job finishes execution will remain until the job is deleted.

- **Job Control object model**
  The Job Control profile provides a mechanism to handle asynchronous execution of methods that take some time to execute.

Job Control object model

The Job Control profile provides a mechanism to handle asynchronous execution of methods that take some time to execute.

For a diagram of how this mechanism works, see Figure 1.

Figure 1. Job Control SMI-S model for IBM FlashSystem A9000 and A9000R systems

For IBMTSDS_MirrorJob, if all the operations completed successfully, the JobState value will be 7 (Completed), and the OperationStatus value will be 2 (Ok), 17 (Completed). If not, the JobState value will be 10 (Exception), and the OperationStatus value will be 6 (Error), 17 (Completed). If the job finished, and the duration time was over 60 minutes, the job will be deleted automatically.

Thin Provisioning profile

Thin provisioning is a capability of some Block Services implementations. It defers provisioning of backing store for regions of a volume until the regions are accessed (written) by the consumer (for example, host file system).

The alternatives (fully provisioned volumes) allocate all the requested capacity from the backing store at the time the volume is created. For thin provisioned volumes, the block server implementation tracks information about which regions are accessed. After a region is accessed, the backing storage is allocated.

The Thin Provisioning profile allows Storage Management Initiative Specification (SMI-S) clients to determine whether a storage system (and children such as pools and volumes) supports thin provisioning. Clients can also determine the difference between the exposed "virtual capacity" and actual committed physical storage and create thinly provisioned volumes and pools.

For IBM FlashSystem A9000 and A9000R, pools are provisioned with virtual capacity. It is not possible to provision a pool with a set amount of physical capacity.

From a volume perspective, you define the volume soft size, and it takes the size from the virtual size of the pool. Upon creation of a volume in a pool, no preallocation of hard capacity is done. Only soft capacity is taken from the virtual capacity of the pool.

- **Thin Provisioning object model**
  Compared with the Block Services profile, the Thin Provisioning profile has more classes and some existing classes with modified properties.
- **Thin Provisioning methods**
  The following sections describe functional methods associated with the Thin Provisioning profile and object model, including uses of and parameters used by each method.

- **Thin Provisioning indications**
  The IBM FlashSystem A9000 and A9000R Common Information Model (CIM) agent supports three Thin Provisioning profile indications.

Thin Provisioning object model

Compared with the Block Services profile, the Thin Provisioning profile has more classes and some existing classes with modified properties.

See Figure 1 for the Thin Provisioning Storage Management Initiative Specification (SMI-S) model for IBM FlashSystem A9000 and A9000R systems.

Figure 1. Thin Provisioning SMI-S model for IBM FlashSystem A9000 and A9000R systems

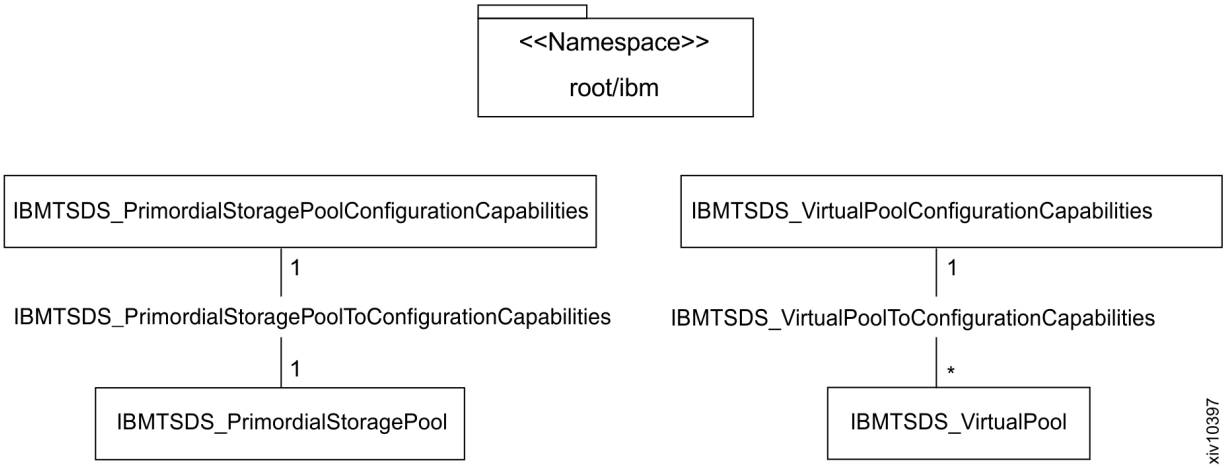

As with the Block Services profile, the CIM agent has four CIM_StorageSetting instances, but IBMTSDS_VirtualPoolSetting adds an instance type in the Thin Provisioning profile. One is the default, and the other is a newly created one.

IBMTSDS_VirtualPoolSetting.InstanceID="IBMTSDS:IBM XIV Virtual Storage Pool Setting"
: This StorageSetting instance is supplied as the Goal parameter for the CreateOrModifyStoragePool method to create or modify a regular thin provisioning VirtualPool instance.

IBMTSDS_VirtualPoolSetting.InstanceID="IBMTSDS:IBM.2810-1310129-112714143137626"
: This StorageSetting instance is supplied as the Goal parameter for the CreateOrModifyStoragePool method to create or modify a thin provisioning VirtualPool instance.

For more information about the thin provisioning object model, see the BlockServices.mof file.

Extrinsic Methods

Table 1. Replication Service methods

Method	Description
CreateSetting	Creates an IBMTSDS_VirtualPoolSetting instance.
modify	Modifies the properties of the newly created IBMTSDS_VirtualPoolSetting instance.
delete	Deletes the newly created IBMTSDS_VirtualPoolSetting instance.
CreateOrModifyStoragePool	Creates or modifies a thin provisioning VirtualPool instance.

Thin Provisioning methods

The following sections describe functional methods associated with the Thin Provisioning profile and object model, including uses of and parameters used by each method.

- **Creating an IBMTSDS_VirtualPoolSetting instance**
  You can use the IBMTSDS_PrimordialStoragePoolCapabilities. CreateSetting method to create an IBMTSDS_VirtualPoolSetting instance. The newly created instance will expire after 600 seconds.
- **Modifying the properties of a newly created IBMTSDS_VirtualPoolSetting instance**
  You can use the IBMTSDS_VirtualPoolSetting. modify method to modify the properties of a newly created IBMTSDS_VirtualPoolSetting instance.
- **Deleting a new IBMTSDS_VirtualPoolSetting instance**
  You can use the IBMTSDS_VirtualPoolSetting. delete method to delete the newly created IBMTSDS_VirtualPoolSetting instance.
- **Creating or modifying a thin provisioning virtual pool**
  You can use the IBMTSDS_StorageConfigurationService. CreateOrModifyStoragePool method to create or modify a thin provisioning VirtualPool instance.
- **Creating a thin provisioning pool**
  You can use the CreateOrModifyStoragePool method to create a thin provisioning pool.

Creating an IBMTSDS_VirtualPoolSetting instance

You can use the IBMTSDS_PrimordialStoragePoolCapabilities. CreateSetting method to create an IBMTSDS_VirtualPoolSetting instance. The newly created instance will expire after 600 seconds.

Parameters

The following list describes the parameters of the IBMTSDS_PrimordialStoragePoolCapabilities. CreateSetting method.

SettingType
: As an input, the value must be 2 (default), 3 (goal), or null.

NewSetting
: As an output, refers to the newly created IBMTSDS_VirtualPoolSetting instance.

Modifying the properties of a newly created IBMTSDS_VirtualPoolSetting instance

You can use the IBMTSDS_VirtualPoolSetting. modify method to modify the properties of a newly created IBMTSDS_VirtualPoolSetting instance.

Parameters

The following list describes the parameters of the IBMTSDS_VirtualPoolSetting. modify method.

ThinProvisionedInitialReserve
: As an input, represents the snapshot size of the newly created pool to be used as the Goal parameter for the CreateOrModifyStoragePool method.

ThinProvisionedPoolType
: As an input, just 7 is supported.

Cop
: As an input, represents the Common Information Model (CIM) object paths (COPs) of the newly created IBMTSDS_VirtualPoolSetting instance to modify.

Deleting a new IBMTSDS_VirtualPoolSetting instance

You can use the IBMTSDS_VirtualPoolSetting. delete method to delete the newly created IBMTSDS_VirtualPoolSetting instance.

Note: The default instance of IBMTSDS_VirtualPoolSetting cannot be deleted.

Parameters

The following list describes the parameters of the IBMTSDS_VirtualPoolSetting. delete method.

Cop
: As an input, represents the Common Information Model (CIM) object paths (COPs) of the newly created IBMTSDS_VirtualPoolSetting instance to delete.

Creating or modifying a thin provisioning virtual pool

You can use the IBMTSDS_StorageConfigurationService. CreateOrModifyStoragePool method to create or modify a thin provisioning VirtualPool instance.

Parameters

The following list describes the parameters of the IBMTSDS_StorageConfigurationService. CreateOrModifyStoragePool method.

Pool *(required for modification)*
: As an input parameter, specifies whether you want to create or modify a pool. If you specify a reference to a pool, indicates that you want to modify the pool. If the parameter is left null, indicates that you want to create a pool.

InPools
: The InPools parameter specifies which pool to create the pool from. To create a VirtualPool instance, only the object reference of IBMTSDS_PrimordialStoragePool is allowed for this parameter since all VirtualPool instances are created in Primordial StoragePool.
To create a SnapshotPool instance, only the object reference of IBMTSDS_VirtualPool is allowed for this parameter, if specified, since a SnapshotPool instance is created in a VirtualPool instance.

  Note: The Common Information Model (CIM) schema defines this input parameter to be an array of strings that represent CIM object paths (COPs), and not actual references to objects.

Goal
: The Goal parameter represents the StorageSetting instance of the pool to be created. To create a VirtualPool instance, only the object reference of IBMTSDS_VirtualPoolSetting is allowed.
To create a SnapshotPool instance, only the object reference of IBMTSDS_SnapshotPoolSetting is allowed. See the `BlockServer.mof` file for details.

ElementName
: The ElementName property provides a means for you to set a meaningful name for the pool that is being created or modified. If specified, it is limited to 63 characters and can contain letters, digits, blank spaces, -, _, . and ~ characters. Blank spaces cannot be the beginning and ending characters. If not specified

during pool creation, a random pool name is generated in the format `pool<random integer>`.

Note: The name of the pool must be unique in the system. It cannot be a name that is already assigned to one of the other pools in the system.

Size *(required for creation)*

As an input parameter, Size specifies the requested size of the pool. Null is not allowed for pool creation. As an output parameter, Size specifies the achieved pool's size.

Creating a thin provisioning pool

You can use the CreateOrModifyStoragePool method to create a thin provisioning pool.

Procedure

1. Call IBMTSDS_PrimordialStoragePool.GetSupportedSizeRange with ElementType set to **`StoragePool`**. If the return code is 0 (success), specify a size.
2. Call invokeMethod on IBMTSDS_PrimordialStoragePoolCapabilities.CreateSetting and specify the SettingType value.
3. Call enumerateInstances on IBMTSDS_VirtualPoolSetting, and then save the CIMObjectPath value of this newly created instance.
4. Call invokeMethod on IBMTSDS_VirtualPoolSetting.modify and specify the IBMTSDS_VirtualPoolSetting value obtained in step 3 to Cop and pool snapshot size from ThinProvisionedInitialReserve.
5. Call invokeMethod on IBMTSDS_StorageConfigurationService.CreateOrModifyStoragePool and specify the IBMTSDS_VirtualPoolSetting value obtained in step 3 to Goal and the ElementName value to Size.
6. Save the pool output parameter. It is the reference to the thin provisioning VirtualPool that was created.

Thin Provisioning indications

The IBM FlashSystem A9000 and A9000R Common Information Model (CIM) agent supports three Thin Provisioning profile indications.

Table 1. Thin Provisioning profile indications

CIM Indication	IBM FlashSystem A9000 and A9000R Event	Description
CAPACITY_WARNING	STORAGE_POOL_VOLUME_USAGE_TOO_HIGH	Thin provisioned volume or pool with the identifier *Volume or Pool ID* capacity in use near available limit.
CAPACITY_CRITICAL	STORAGE_POOL_EXHAUSTED	Thin provisioned volume or pool with the identifier *Volume or Pool ID* capacity in use exceeded available limit.
CAPACITY_OKAY	STORAGE_POOL_VOLUME_USAGE_BACK_TO_NORMAL	Thin provisioned volume or pool with the identifier *Volume or Pool ID* capacity condition cleared.

Conformance tests

All Storage Networking Industry Association (SNIA) official Certification Test Programs (CTP) and Microsoft System Center Virtual Machine Manager (SCVMM) Storage Automation tests have been completed.

SNIA official CTP tests

Table 1. SNIA CTP results

IBM FlashSystem A9000 and A9000R CIM agent release	Official CTP Test Suite	Official CTP Test Results
12.0.x	• Storage Management Initiative Specification (SMI-S) version 1.6 • Test version 1.6.0.1081	http://www.snia.org/tech_activities/standards/curr_standards/smi

SCVMM Storage Automation tests

All Microsoft System Center Virtual Machine Manager (SCVMM) 2012 Storage Automation tests were completed on IBM FlashSystem A9000 and A9000R with IBM FlashSystem A9000 and A9000R Common Information Model (CIM) agent version 12.0.

Remote support

Refer to the following sections for information about remote support available for IBM FlashSystem® A9000R.

- **Remote support and proactive support**
  To allow IBM® to provide support for the storage system, the proactive support and remote support options are available.

- **Remote Support Proxy**
  Refer to the following sections for information about installing, configuring, and managing IBM Remote Support Proxy.

Remote support and proactive support

To allow IBM® to provide support for the storage system, the proactive support and remote support options are available.

Note: For various preventive and diagnostics support actions relating to the storage system's continuous operation, IBM Support requires customer approval. Without customer approval, these support actions cannot be performed.

- **Proactive support** ("Call Home") – Allows proactive notifications regarding the storage system health and components to be sent to IBM Support at predefined intervals. Heartbeats and events are sent from the system to the IBM service center. The service center analyzes the information within the heartbeats and the events, correlates it with its vast database and can then trigger a component replacement prior to its potential failure.
  Upon detection of any hardware or software error code, both IBM Support and your predefined contact person are notified via e-mail, through a specified SMTP gateway. If IBM Support determines that the detected event requires service or further investigation, a new PMR is created and sent to appropriate IBM Support team. Proactive support minimizes the number of interaction cycles with IBM Support.
  Note: Proactive support is compliant with General Data Protection Regulation (GDPR) EU 2016/679. For more information, refer to http://ibm.biz/gdpr-storage-paper.
- If required, the customer email gateway can be configured to send call home information to IBM only via a secured channel. For more information, refer to the 'Encrypting Call Home and heartbeat notifications' section in the IBM FlashSystem A9000 and IBM FlashSystem A9000R Architecture and Implementation Redbooks publication.

- **Remote support** – Allows IBM Support to remotely and securely access your storage system when needed during a support call. This option requires IP communication between the storage system and the IBM Remote Support Center. If a storage system does not have direct access to the Internet (for example, due to a firewall), use the IBM Remote Support Proxy utility to enable the connection. Remote support minimizes the time it takes to diagnose and remedy storage system operational issues.

For more information, refer to the *Planning for remote support, on-site service, and maintenance* section of the *Deployment Guide*.

Remote Support Proxy

Refer to the following sections for information about installing, configuring, and managing IBM® Remote Support Proxy.

Introduction

The remote support system for IBM® storage systems requires TCP/IP communication between the storage system and the IBM Remote Support Center. When a storage system does not have direct access to the Internet (for example, due to a firewall), you can use the Remote Support Proxy to facilitate that connection.

The IBM® Remote Support Proxy utility creates a network proxy that connects one or more IBM storage systems to IBM remote-support servers in the IBM Remote Support Center. It establishes a service on a Linux® system that has Internet connectivity to the IBM Remote Support Center and local network connectivity to the storage system. The connection to the IBM Remote Support Center is initiated by the storage system through its management graphical user interface (GUI) or command-line interface (CLI).

The figure below illustrates a typical network configuration that uses the Remote Support Proxy. The communication between the storage system and the Remote Support Proxy uses Secure Shell (SSH). The communication between the Remote Support Proxy and the Remote Support Center is encrypted with an additional layer of Secure Sockets Layer (SSL).

Figure 1. Remote Support Proxy components

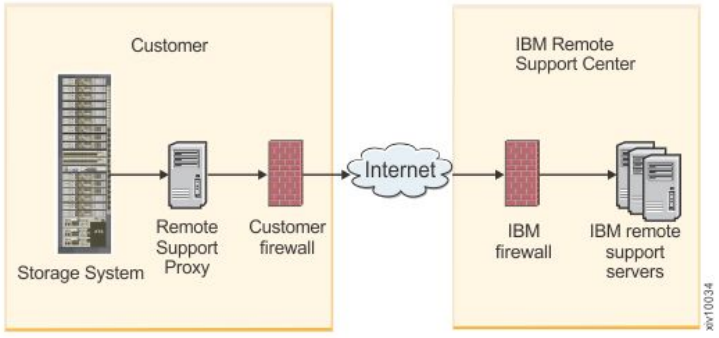

- **Compatibility and requirements**
  For the complete and up-to-date information about the compatibility and requirements of the IBM Remote Support Proxy utility, refer to its latest release notes.

Compatibility and requirements

For the complete and up-to-date information about the compatibility and requirements of IBM® Remote Support Proxy utility, refer to its latest release notes.

Installation

Use this information to either install or upgrade the IBM® Remote Support Proxy utility.

- **Installing and configuring for the first time**
  This information describes how to install and configure the IBM Remote Support Proxy utility on a system that has access to an IBM storage system and the Internet.
- **Upgrading and configuring an existing installation**
  This information describes how to upgrade to a newer version of the IBM Remote Support Proxy utility on a system that has access to a storage system and the Internet.
- **Uninstalling the Remote Support Proxy**
  This information describes how to uninstall the IBM Remote Support Proxy utility.

Installing and configuring for the first time

This information describes how to install and configure the IBM® Remote Support Proxy utility on a system that has access to an IBM storage system and the Internet.

- **Installing the software package**
  This information describes how to download and install the IBM Remote Support Proxy utility.
- **Configuring the Remote Support Proxy service**
  Configure the Remote Support Proxy service before starting it.
- **Starting the Remote Support Proxy service**
  After updating the Remote Support Proxy configuration file, start the Remote Support Proxy service.
- **Configuring storage systems**
  After installing and configuring the Remote Support Proxy, define a new remote support connection with information about the proxy server.

Installing the software package

This information describes how to download and install the IBM® Remote Support Proxy utility.

Procedure

To download and install the IBM Remote Support Proxy utility, perform these steps:

1. Download the Remote Support Proxy utility installation package and *Release Notes* from IBM Fix Central.
2. Assign the downloaded file execution permissions by running the following command:

   ```
   chmod +x file_name.bin
   ```

   Where *file_name*.bin is the name of the installation package for the appropriate platform.

3. Install the IBM Remote Support Proxy utility using the following command:

   ```
   ./file_name.bin
   ```

   Where *file_name*.bin is the name of the installation package for the appropriate platform.

   A software license is displayed.

4. Follow the instructions to read through and accept the license.
   The installation program installs an RPM package named xivproxy and then exits.

Results

The following files are installed on the local system:

/usr/bin/xivproxy
: The binary executable file.

/etc/xiv/proxy.conf
: The configuration file. You must update this file before starting the Remote Support Proxy service.

/etc/init.d/xivproxy
: The system service for starting and stopping the Remote Support Proxy. The service is configured to start when the server boots. It does not start automatically after installation.

/usr/share/xiv/syslog-logger
: The logger program that can be modified to integrate with the existing logging and monitoring systems.

/usr/share/doc/xivproxy-*version*/license
: The directory containing the accepted software license for the Remote Support Proxy in multiple languages.

In the directory path, *version* is the product version.

/usr/share/xiv/proxy_id_rsa
: A cryptographic file used by the configuration-retrieval function.
Important: If a restricted user is specified by using the User configuration parameter, that user must be able to read this file.

/var/lib/xiv/proxy
: The directory used to store configuration updates.
Important: If a restricted user is specified by using the User configuration parameter, that user must be able to write to this directory.

Configuring the Remote Support Proxy service

Configure the Remote Support Proxy service before starting it.

Procedure

To configure the Remote Support Proxy service, edit the following parameters in the /etc/xiv/proxy.conf file:

ListenInterface
: The IPv4 address, IPv6 address, or interface name on which the proxy service is to listen for incoming connections from the storage system (for example, `192.0.2.1`, `2001:DB8:0:0:0:0:0:0`, or `eth0`).
If you specify an IPv4 or IPv6 address, the proxy service listens on only the specified IP address.

If you specify an interface name, the proxy service listens on all allowed addresses. You can configure which addresses are allowed by setting the UseIPv4, UseIPv6, and UseIPv6LinkLocalAddress parameters. By default, the proxy service listens on both IPv4 and IPv6 addresses.

If you specify the interface name, the name must be the same as the output from the ifconfig command.

Important: The storage system must have access to the system running the Remote Support Proxy as defined by this interface name or IP address.

ListenPort
: The TCP port on which the Remote Support Proxy is to listen for incoming connections from the storage system. For example, to listen on port 9955, add the following strings to the configuration file:

```
# Interface and port on which the proxy listens
ListenInterface eth0
ListenPort 9955
```

Note: If ListenPort is not specified, the default port used is `8988`.
Important: The storage system must have access to the system running the Remote Support Proxy as defined by this port.

UseIPv6
: (Optional) When an interface name is specified for the ListenInterface or StatusInterface parameter, this parameter determines whether the proxy service listens on the interfaces IPv6 addresses (for example, `2001:DB8:0:0:0:0:0:0`) if they exists.
This parameter has no effect if an interface does not have an IPv6 address.

If UseIPv6 is set to `yes`, the proxy service listens on the IPv6 address if one exists. This is the default value.

To also listen on an IPv6 link-local address, UseIPv6LinkLocalAddress must be set to `yes`.

If UseIPv6 is set to `no`, the proxy service does not listen on any IPv6 address, including the IPv6 link-local address. In this case, the interface must have an IPv4 address defined, and UseIPv4 cannot be set to `no`.

Important: Setting UseIPv6 to `no` does not affect outgoing communication from the Remote Support Proxy to the Remote Support Center. This communication might still use the IPv6 protocol.

UseIPv6LinkLocalAddress
: (Optional) When an interface name is specified for the ListenInterface or StatusInterface parameter, this parameter determines whether the proxy server listens on the IPv6 link-local address (for example, `FE80:DB8:0:0:0:0:0:0`).
By default, UseIPv6LinkLocalAddress is set to `no`.

Notes:

- If a numeric IPv6 link-local address is used in ListenInterface or StatusInterface, this parameter must be set to `yes`.
- UseIPv6LinkLocalAddress must not be set `yes` when UseIPv6 is set to `no`.

UseIPv4
: (Optional) When an interface name is specified for the ListenInterface or StatusInterface parameter, this parameter determines whether the proxy service listens on the interface's IPv4 addresses (for example, `192.0.2.1`) if they exists.
This parameter has no effect if an interface does not have an IPv4 address.

If UseIPv4 is set to `yes`, the proxy service listens on the IPv4 address if one exists. This is the default value.

If UseIPv4 is set to `no`, the proxy service does not listen on the IPv4 address. In this case, the interface must have an IPv6 address defined, and UseIPv6 cannot be set to `no`.

Important: Setting UseIPv6 to `no` does not affect outgoing communication from the Remote Support Proxy to the Remote Support Center. This communication might still use the IPv6 protocol.

HTTPProxyHost
: (Optional) The IPv4 address, IPv6 address, or host name of the HTTP proxy server (for example, `192.0.2.3`).
If your network does not allow a direct connection to the Internet, you must use an HTTP proxy server. If so, set the HTTPProxyHost and HTTPProxyPort parameters.

Note: The HTTP proxy must be configured to allow outgoing HTTPS connections.

HTTPProxyPort

(Optional) The port of the HTTP proxy server (for example, `8080`).

HTTPProxyUser

(Optional) The user ID used to authenticate to the HTTP proxy (for example, `ibm_user`).

By default, the Remote Support Proxy does not use HTTP authentication. If authentication is required to access sites outside your network, set the HTTPProxyUser and HTTPProxyPassword parameters.

Note: The only supported method is basic authentication. Challenge/Response methods, such as NTLM, are not supported.

HTTPProxyPassword

(Optional) The password used to authenticate to the HTTP proxy (for example, `passw0rd`).

StatusInterface

(Optional) The IPv4 address, IPv6 address, or interface name of an additional TCP port on which the proxy service listens to monitor status (for example, `192.0.2.1`, `2001:DB8:0:0:0:0:0:0`, or `eth0`).

If you specify an IPv4 or IPv6 address, the proxy service listens on only the specified IP address.

If you specify an interface name, the proxy service listens on all allowed addresses. You can configure which addresses are allowed by setting the UseIPv4, UseIPv6, and UseIPv6LinkLocalAddress parameters. By default, the proxy service listens on both IPv4 and IPv6 addresses.

When you connect to the status port (for example, by using **telnet**), the proxy service displays a status message and then closes the connection.

If the StatusInterface and StatusPort parameters are not configured, the status service is disabled.

When the status service is enabled, you can retrieve the status by running the service xivproxy status command.

To enable status queries only from the local server, use the loopback interface `lo`. To enable status queries over the network, use a network interface.

StatusPort

(Optional) The port number of an additional TCP port on which the Remote Support Proxy listens to monitor status (for example, `8966`).

If the StatusInterface and StatusPort parameters are not configured, the status service is disabled.

User

(Optional) The user name of a restricted user that is used to run the Remote Support Proxy (for example, `nobody`).

If the proxy runs as the root user, it drops its privileges to match the privileges of the specified user or to `nobody` if no user is specified. This is done after binding ports and opening the log file, in case they require root access.

Logger

(Optional) Location of the logger program (for example, /usr/share/xiv/syslog-logger).

See the contents of the /usr/share/xiv/syslog-logger file for more information and example code.

Note: The logger runs as a restricted user and not the root user.

The Remote Support Proxy runs the logger each time the storage system connects to and disconnects from the Remote Support Center, or when an error occurs. You can modify the logger program to integrate with existing logging and monitoring tools.

The logger program runs without any arguments and with the following predefined environment variables:

```
PATH=/bin:/sbin:/usr/bin:/usr/sbin
EVENT={connect|disconnect|error}
CLIENT=IP_address:port
SERVER=IP_address:port
SERVER_LIST=IP_address:port,IP_address:port,...
MESSAGE=arbitrary text message when EVENT=error
```

SERVER_LIST is a list of all known IBM® remote-support servers at the time of the last configuration retrieval.

In the following example, the storage system uses IP address 192.0.2.5 on port 9999, the Remote Support Center uses the following IP addresses on port 443 and the logger is called to report a successful connection:

- 129.33.206.139
- 204.146.30.139
- 170.225.126.11
- 170.225.126.12
- 170.225.127.11
- 170.225.127.12

```
PATH=/bin:/sbin:/usr/bin:/usr/sbin
EVENT=connect
CLIENT=192.0.2.5:9999
SERVER=129.33.206.139:443
SERVER_LIST=129.33.206.139:443,204.146.30.139:443,170.225.126.11:443,170.225.126.12:443,170.225.127.11:443,170.225.1
27.12:443
MESSAGE=
```

LogFile

(Optional) Location of the error log. If not set, the default location is /var/log/xivproxy.log.

This error log contains information about proxy service startup, shutdown, connections, and disconnections. If debugging is enabled, this file also contains the trace logs.

DebugLog

(Optional) A flag indicating whether debugging is enabled. Set to `yes` to enable debugging or `No` to disable debugging. The default value is `no`.

When debugging is enabled, trace logs are created. These logs are useful for troubleshooting and debugging problems.

ServerAddress

(Optional) IP addresses of front servers. If set, only the configured front server(s) will be used. Below is the list of IP addresses of all the available front servers:

- 129.33.206.139
- 204.146.30.139
- 170.225.126.11
- 170.225.126.12
- 170.225.127.11
- 170.225.127.12

To configure front servers, use the following format:

```
ServerAddress[N] IPAddrs
ServerPort[N]PortNo (where N=1..n)
```

For example, to configure two front servers with IP addresses IP1 and IP2 over SSL, add the following to /etc/xiv/proxy.conf:

```
ServerAddress1 IP1
ServerPort1 8888
ServerAddress2 IP2
ServerPort2 8888
```

Note: If ServerPort is not specified, the default port used is `443`.

Example

The following example shows a sample configuration file:

```
# Interface and port on which the proxy listens
ListenInterface          eth0
ListenPort               8988

# IP address settings
UseIPv4                  yes
UseIPv6                  yes
UseIPv6LinkLocalAddress  no

# Host and port of the HTTP proxy server
HTTPProxyHost            192.0.2.3
HTTPProxyPort            8080

# User ID and password for the HTTP proxy server
HTTPProxyUser            ibm_user
HTTPProxyPassword        passw0rd

# Interface and port of status service
StatusInterface          lo
StatusPort               8966

# User ID to run the proxy
User                     nobody

# External log program (optional)
Logger                   /usr/share/xiv/syslog-logger
LogFile                  /var/log/xivproxy.log
DebugLog                 no

# Front Server IP address and port (optional)
ServerAddress1
IP1ServerPort1           8888
```

Starting the Remote Support Proxy service

After updating the Remote Support Proxy configuration file, start the Remote Support Proxy service.

Procedure

To start the Remote Support Proxy service, run the following command:

```
service xivproxy start
```

If the proxy service does not start correctly, examine the log file for errors. The default log file is /var/log/xivproxy.log.

Configuring storage systems

After installing and configuring the Remote Support Proxy, define a new remote support connection with information about the proxy server.

Procedure

1. Ensure that no remote support servers are configured. To list all currently configured Remote Support Center servers, use the following command:

   ```
   support_center_list
   ```

2. If Remote Support Center servers have been defined, delete the servers by using the following command:

   ```
   support_center_delete support_center=name
   ```

   Where name is the unique name of the Remote Support Center server, as it appears in the Names column in the output of the support_center_list command.

3. Define a new remote support connection with the information about the proxy server by using the following command:

   ```
   support_center_define support_center=name address=proxy_ip port=proxy_port
   ```

   where:

   support_center=name
   : A unique name of the remote support server. You can use any name.

   address=proxy_ip
   : The IP address, in IPv4 or IPv6 format, of the server on which the Remote Support Proxy is running. The value is the same as the ListenInterface parameter in the Remote Support Proxy configuration file (/etc/xiv/proxy.conf).

   port=proxy_port
   : The TCP port on which the Remote Support Proxy is to listen for incoming connections from the storage system (for example, `8988`). The value is the same as the ListenPort parameter in the Remote Support Proxy configuration file (/etc/xiv/proxy.conf).

4. Test the connection for each storage system when configuration is complete by connecting the storage system to the remote support server by using the following command:

   ```
   support_center_connect
   ```

5. Wait a few seconds for the connection to complete, and then check the status by using the following command:

   ```
   support_center_status
   ```

6. Ensure that the connection state is reported as **`idle`** and not as **`no connection`**.
   If the connection state is not **`idle`**, examine the log file for errors.

Upgrading and configuring an existing installation

This information describes how to upgrade to a newer version of the IBM Remote Support Proxy utility on a system that has access to a storage system and the Internet.

Before you begin

If the Remote Support Proxy service is running when you start the installation, the service is stopped when the installation completes.

Procedure

Perform these steps to upgrade and configure the IBM Remote Support Proxy utility:

1. Download and install the Remote Support Proxy (see Installing the software package).
   The installer recognizes the existing installation and performs an upgrade.

2. Configure the Remote Support Proxy (see Configuring the Remote Support Proxy service).
   Review the configuration carefully to ensure that the Remote Support Proxy is correctly configured after upgrading. All configuration settings are listed in Configuring the Remote Support Proxy service.

3. Start the Remote Support Proxy by running the following command:

   ```
   service xivproxy start
   ```

   If the proxy service does not start correctly, examine the log file for errors. The default log file is /var/log/xivproxy.log.

Uninstalling the Remote Support Proxy

This information describes how to uninstall the IBM® Remote Support Proxy utility.

Procedure

To remove the IBM Remote Support Proxy utility from the system on which it is installed, perform these steps:

1. Stop the proxy service running the following command:

   ```
   service xivproxy stop
   ```

2. Uninstall the IBM Remote Support Proxy utility by running the following command:

```
rpm -e xivproxy
```

Management

This information describes how to start and stop the Remote Support Proxy and how to view its connection status.

- **Starting and stopping the Remote Support Proxy service**
  After the IBM Remote Support Proxy utility is installed and configured, you can start, stop, and restart the Remote Support Proxy service.
- **Monitoring status**
  You can monitor connection status of the Remote Support Proxy.
- **Updating remote-support server information**
  The Remote Support Proxy periodically connects to the Remote Support Center to retrieve configuration updates, including the current IP address and port for each IBM® remote-support servers.

Starting and stopping the Remote Support Proxy service

After the IBM Remote Support Proxy utility is installed and configured, you can start, stop, and restart the Remote Support Proxy service.

Procedure

- To start the Remote Support Proxy, run the following command:

  ```
  service xivproxy start
  ```

- To stop the Remote Support Proxy, run the following command:

  ```
  service xivproxy stop
  ```

- To restart the Remote Support Proxy, run the following command:

  ```
  service xivproxy restart
  ```

Monitoring status

You can monitor connection status of the Remote Support Proxy.

Procedure

To monitor Remote Support Proxy status, run the following command:

```
service xivproxy status
```

Results

The status command lists the following values.
Note: This information is displayed only if the status service is configured in the configuration file.

Open connections
: The number of connections that currently exist between the storage system on the network and the IBM® Remote Support Center server. These connections do not have to be in an active remote support session, but only connected to the server.

Failed connections
: The number of connections that failed due to an error. This number represent a cumulative count from the time the status service was started and can only increase. When the proxy service is stopped and restarted, this number resets to zero.

Total connections
: The total number of connection attempts that were made by the storage system on the network, whether successful or not. This number is a cumulative count from the time the proxy service was started and can only increase. When the proxy service is stopped and restarted, this number resets to zero.

Updating remote-support server information

The Remote Support Proxy periodically connects to the Remote Support Center to retrieve configuration updates, including the current IP address and port for each IBM® remote-support servers.

About this task

No user action is required to perform this update. This retrieval is performed automatically by the Remote Support Proxy every six hours. The connection is made by using the same method (direct connection or through an HTTP proxy) that is used to create a remote-support connection to the storage system.

The Remote Support Proxy connects directly to the Remote Support Center with no involvement from the storage system.

The configuration updates are saved to a file in the /var/lib/xiv/proxy directory.

Publications and related information

The following Portable Document Format (PDF) publications and related online information sources are available for IBM FlashSystem® A9000R 12.3.x.

To view a PDF file, you need Adobe Reader. You can download it at no charge from the Adobe website.

Note:

- For release note publications, see Release notes.
- To access PDF versions of publications, see the Publications and related information within the IBM Hyper-Scale Manager documentation.

Publications

IBM FlashSystem A9000R publication	Description	PDF download link
Product Overview	This publication provides introductory and conceptual information about the IBM FlashSystem A9000R functional features. Second Edition (August 2022)	Download
Deployment Guide	This publication describes the hardware deployment and other preparation tasks for IBM FlashSystem A9000R. First Edition (March 2022)	Download
CLI Reference Guide (version 12.3.2.e)	This publication describes the available command-line interface (CLI) commands, parameter definitions, and return codes for IBM FlashSystem A9000R 12.3.2.e. First Edition (January 2023) Note: The publication released for version 12.3.2.e is also applicable to versions 12.3.2.a, 12.3.2.b, 12.3.2.c, and 12.3.2.d	Download
CLI Reference Guide (12.3.0 - 12.3.2)	This publication describes the available command-line interface (CLI) commands, parameter definitions, and return codes for IBM FlashSystem A9000R for versions 12.3.0 - 12.3.2. Publication number: SC27-8711-12	Download
Open API Reference Guide	This publication describes the IBM FlashSystem A9000 and IBM FlashSystem A9000R Open Application Programming Interface (API). Publication number: SC27-8561-04 Note: The publication released for version 12.0.3 is also applicable to version 12.3.1.	Download

IBM Spectrum Accelerate Family HyperSwap Quorum Witness publication	Description	PDF download link
User Guide (version 1.0.3)	This publication describes how to install and operate IBM Spectrum Accelerate Family HyperSwap Quorum Witness, a software component required for IBM HyperSwap solution. Publication number: SC27-4631-10 Note: This publication is relevant for version 1.0.3 only. Use this version if you are running Red Hat® Enterprise Linux® (RHEL) or CentOS version 8.2.	Download
User Guide (version 1.0.2)	This publication describes how to install and operate IBM Spectrum Accelerate Family HyperSwap Quorum Witness, a software component required for IBM HyperSwap solution. Publication number: SC27-4631-08 Note: This publication is relevant for version 1.0.2, for operating systems running Red Hat Enterprise Linux (RHEL) and CentOS *lower than* version 8.x.	Download

XIV® Remote Support Proxy publication	Description	PDF download link
Installation and User Guide	This publication provides installation and user guidance for IBM Remote Support Proxy version 1.3.x. Ninth Edition (July 2022)	Download

Customer self-support information

To IBM customers who are not receiving on-site service from IBM:

If you wish to obtain a license for an IBM proprietary software tool that facilitates the maintenance of IBM FlashSystem A9000, FlashSystem A9000R, or XIV Gen3 systems (the *Support Tool*), send an email request to the following email address: askaccel@us.ibm.com.

The terms of this license will allow you, your agent or contractor to use the Support Tool solely to provide maintenance for your IBM storage product. You remain responsible for ensuring that any such agent or contractor is in compliance with the terms and conditions of the IBM license. Please note that use of the Support Tool is provided on an 'as is' basis and should be used only by trained, skilled personnel. IBM is not responsible for issues arising from use of the Support Tool and reserves the right to assist you in resolving such issues on a time and materials basis.

Related information

On the following online information sources you can find additional information and publications related to IBM FlashSystem A9000R.

IBM information source	Description
IBM IT Infrastructure webpage	Data sheets, white papers, videos, and other marketing materials for IBM Flash Storage products
IBM FlashSystem A9000 documentation	Product documentation for IBM FlashSystem A9000
IBM Storage Redbooks®	IBM Redbooks and Redpapers on IBM storage products

External information source	Description	PDF download link
FSTC RFID IT Assets – Phase 2 Business Case and Best Practices	This report contains RFID requirements and specifications, published by the Financial Services Technology Consortium (FSTC).	Download

Glossary

This glossary provides terms and definitions for FlashSystem software and products.

The following cross-references are used in this glossary:

- *See* refers you from a nonpreferred term to the preferred term or from an abbreviation to the spelled-out form.
- *See also* refers you to a related or contrasting term.

For other terms and definitions, see the IBM® Terminology website (opens in new window).

A B C D E F G H I J L M N O P Q R S T U V W Z

A

access mode
: One of the modes in which a logical unit (LU) in a disk controller system can operate. The three access modes are image mode, managed space mode, and unconfigured mode. See also image mode, managed mode, unconfigured mode.

Address Resolution Protocol (ARP)
: A protocol that dynamically maps an IP address to a network adapter address in a local area network.

advisory lock
: A type of lock that a process holds on a region of a file that signals any other process to not use or lock the region or an overlapping region. Other processes are not forced to comply.

agent code
: An open-systems standard that interprets Common Information Model (CIM) requests and responses as they are transferred between the client application and the device.

alert
: An informational event that requires user action. See also event.

application server
: A host that is attached to the storage area network (SAN) and that runs applications.

ARP
: See Address Resolution Protocol.

array
: 1. An ordered collection, or group, of physical devices (for example, drives), that provide storage for virtual volumes. An array is a group of drives designated to be managed with a Redundant Array of Independent Disks (RAID). See also Redundant Array of Independent Disks.
  2. An ordered collection, or group, of physical devices (disk drive modules) that are used to define logical volumes or devices. An array is a group of drives designated to be managed with a Redundant Array of Independent Disks (RAID).

association
: A class that contains two references that define a relationship between two referenced objects.

asymmetric virtualization
: A virtualization technique in which the virtualization engine is outside the data path and performs a metadata-style service. The metadata server contains all the mapping and locking tables while the storage devices contain only data. See also symmetric virtualization.

asynchronous replication
: 1. A type of replication in which control is given back to the application as soon as the write operation is made to the source volume. Some time later, the write operation is made to the target volume. See also synchronous replication.
  2. In replication, the process of copying data from a source table to a target table outside the scope of the original transaction that updated the source table. See also synchronous replication.

audit log
: An unalterable record of all commands or user interactions that are issued to the system.

authenticated user
: A user who has logged in to the system with a valid account (user ID and password).

auxiliary volume
: The volume that contains a mirror of the data on the master volume. See also master volume, relationship.

availability
: The delivery of reliable service during scheduled periods.

B

backup
: A copy of a data set or object to be used in case of accidental loss.

bandwidth
: The amount of information that a system can transfer in a given time.

bitmap
: A coded representation in which each bit, or group of bits, represents or corresponds to an item; for example, a configuration of bits in main storage in which each bit indicates whether a peripheral device or a storage block is available or in which each group of bits corresponds to one pixel of a display image.

blade
: See blade server.

blade server
: An independent server containing one or more processors, memory, disk storage, and network controllers. A blade server runs its own operating system and applications.

block
: 1. A unit of data storage on a device.
  2. To suspend a program process.

block storage
: A unit of data storage on a device.

Boolean
: Characteristic of an expression or variable that can only have a value of true or false.

boot drive
: A drive that has the required software to start a system.

C

cache
: 1. Memory used to improve access times to instructions, data, or both. Data that resides in cache memory is normally a copy of data that resides elsewhere in slower, less expensive storage, such as on a disk or on another network node.
  2. Storage used to improve access times to instructions, data, or both. For example, data that resides in cache memory is normally a copy of data that resides elsewhere in slower, less expensive storage, such as on a disk or on another network node.

cache eviction
: A process by which data associated with a file is removed from the cache system. The data is removed either by using a Least Recently Used (LRU) algorithm when configured General Parallel File System (GPFS) hard or soft quota limits are exceeded or by issuing a command. When referenced again in the cache system, the data that is associated with the file is retrieved from the home system.

caching I/O group
: The I/O group in the system that performs the cache function for a volume.

call home
: A communication link established between a product and a service provider. The product can use this link to place a call to IBM or to another service provider when it requires service. With access to the machine, service personnel can perform service tasks, such as viewing error and problem logs or initiating trace and dump retrievals.

canister
: A single processing unit within a storage system.

capacity licensing
: A licensing model that licenses features with a price-per-terabyte model. Licensed features are FlashCopy®, and Metro Mirror and Global Mirror, and virtualization. See also FlashCopy, Metro Mirror, virtualization.

chain
: A set of enclosures that are attached to provide redundant access to the drives inside the enclosures. Each control enclosure can have one or more chains.

Challenge Handshake Authentication Protocol (CHAP)
: An authentication protocol that protects against eavesdropping by encrypting the user name and password.

change volume
: A volume that is used in Global Mirror that holds earlier consistent revisions of data when changes are made.

CHAP
: See Challenge Handshake Authentication Protocol.

CHAP secret
: In the Challenge-Handshake Authentication Protocol (CHAP), a secret passphrase that is used to authenticate a storage system to iSCSI-attached hosts.

CIFS
: See Common Internet File System.

CIM
: See Common Information Model.

CIM object manager (CIMOM)
: The common conceptual framework for data management that receives, validates, and authenticates the CIM requests from the client application. It then directs the requests to the appropriate component or service provider.

CIMOM
: See CIM object manager.

class
: The definition of an object within a specific hierarchy. A class can have properties and methods and can serve as the target of an association.

CLI
: See command-line interface.

client
: A software program or computer that requests services from a server. See also host.

client application
: A storage management program that initiates Common Information Model (CIM) requests to the CIM agent for the device.

clone
: A copy of a volume on a server at a particular point in time. The contents of the copy can be customized while the contents of the original volume are preserved.

cluster

1. See clustered system.
2. A group of computers and other resources that operate together as a single system. See also clustered system, GPFS cluster.

clustered system
: A collection of nodes that are placed in pairs (I/O groups) for redundancy, which provide a single management interface. See also cluster, GPFS cluster, system.

command-line interface (CLI)
: A computer interface in which the input and output are text based.

Common Information Model (CIM)
: An implementation-neutral, object-oriented schema for describing network management or systems management information. The Distributed Management Task Force (DMTF) develops and maintains CIM specifications.

Common Internet File System (CIFS)

1. A protocol that manages shared, remote file access for applications to files, printers, serial ports, and so on over a TCP/IP network.
2. A protocol that enables collaboration on the Internet by defining a remote file-access protocol that is compatible with the way applications already share data on local disks and network file servers. See also Server Message Block 2.0, Server Message Block.

compression
: A function that removes repetitive characters, spaces, strings of characters, or binary data from the data being processed and replaces characters with control characters. Compression reduces the amount of storage space that is required for data.

compression accelerator
: Hardware onto which the work of compression is off-loaded from the microprocessor.

compute node
: An independent server that contains one or more microprocessors, memory, storage, and network controllers and runs its own operating system and applications.

concurrent copy
: A function of the DFSMSdss component that is used to back up any collection of data at a point in time with minimum down time for the database or application that uses the collection of data.

configuration node
: A node that acts as the focal point for configuration commands and manages the data that describes the clustered-system configuration.

consistency group
: A group of copy relationships between virtual volumes or data sets that are maintained with the same time reference so that all copies are consistent in time.

container
: A software object that holds or organizes other software objects or entities.

contingency capacity
: For thin-provisioned volumes that are configured to automatically expand, the unused real capacity that is maintained. For thin-provisioned volumes that are not configured to automatically expand, the difference between the used capacity and the new real capacity.

control enclosure
: A hardware unit that includes the enclosure chassis, node canisters, drives, and system function.

controller
: See node canister.

controller unit
: See control enclosure.

counterpart SAN
: A non-redundant portion of a redundant storage area network (SAN). A counterpart SAN provides all the connectivity of the redundant SAN but without the redundancy. Each counterpart SAN provides an alternate path for each SAN-attached device. See also redundant SAN.

count key data

1. In System z®, a data recording format that uses self-defining record formats in which each record on a volume is represented by up to three fields: a count field identifying the record and specifying its format, an optional key field that can be used to identify the data area contents, and an optional data field that typically contains the user data. See also storage architecture type.
2. An architecture for a direct access storage device (DASD) device or logical device that specifies the access mechanisms for the logical data units on the device through a specific set of supported channel commands. Extensions to the CKD command set form the basis of Extended CKD.

cross-volume consistency
: A consistency group property that guarantees consistency between volumes when an application issues dependent write operations that span multiple volumes.

CRU
: See customer-replaceable unit.

customer-replaceable unit (CRU)
: An assembly or part that a customer can replace.

cylinder
: A unit of storage on a count-key-data (CKD) device with a fixed number of tracks.

D

data consistency
: A characteristic of the data at the target site where dependent write order is maintained to guarantee the recoverability of applications.

data migration
: The movement of data from one physical location to another without disruption of application I/O operations.

data record

1. A basic unit of data recording format that is also known as a count-key-data (CKD) record.
2. The basic unit of S/390® and zSeries storage on an ESS, also known as a count-key-data (CKD) record. See also count key data, fixed-block architecture.

dense wavelength division multiplexing (DWDM)

A technology that places many optical signals onto one single-mode fiber using slightly different optical frequencies. DWDM enables many data streams to be transferred in parallel.

dependent write operation
A write operation that must be applied in the correct order to maintain cross-volume consistency.

destage
To move data from cache to a nonvolatile storage medium.

device
A piece of equipment such as a workstation, printer, disk drive, tape unit, or remote system.

device provider
A device-specific handler that serves as a plug-in for the Common Information Model (CIM); that is, the CIM Object Manager (CIMOM) uses the handler to interface with the device.

DIMM
See dual inline memory module.

directed maintenance procedure
See fix procedure.

discovery
The automatic detection of a network topology change, for example, new and deleted nodes or links.

disk
See drive.

disk drive
See drive.

disk drive bay
See drive.

disk drive module
See drive.

disk zone
A zone that is defined in the storage area network (SAN) fabric in which the systems can detect and address the logical units (LUs) that the storage systems present.

Distributed Management Task Force (DMTF)
An alliance of computer vendors that was convened to define streamlined management of the diverse operating systems commonly found in an enterprise.

DMTF
See Distributed Management Task Force.

DNS
See domain name server.

domain name server (DNS)
A server program that supplies name-to-address conversion by mapping domain names to IP addresses.

DRAM
See dynamic random access memory.

drive
A data storage device. A drive can be either a magnetic disk drive or a solid-state drive (SSD).

DS Storage Manager
See graphical user interface.

dual inline memory module (DIMM)
A small circuit board with memory-integrated circuits containing signal and power pins on both sides of the board.

DWDM
See dense wavelength division multiplexing.

dynamic random access memory (DRAM)
Storage in which the cells require repetitive application of control signals to retain stored data.

E

EC
See engineering change.

ECKD
See extended count key data.

EIA
See Electronic Industries Alliance.

Electronic Industries Alliance (EIA)
An organization of electronics manufacturers that advances the technological growth of the industry, represents the views of its members, and develops industry standards.

enclosure

1. The metal structure in which various electronic components are mounted.
2. A unit that houses the components of a storage subsystem, such as a control unit, disk drives, and power source.

end cap
A removable hardware piece that fits over the end of each side of a drive assembly.

engineering change (EC)
An update to a machine, part, or program. Each EC for a given unit is assigned a unique number referred to as an EC level or EC number.

environment service module canister
See expansion canister.

error
See event.

EUI
See extended unique identifier.

event

An occurrence of significance to a task or system. Events can include completion or failure of an operation, a user action, or the change in state of a process. See also alert, message.

exclude
: To prevent the use of a managed disk (MDisk) by a clustered system because of certain error conditions.

expansion canister
: A hardware unit that includes the serial-attached SCSI (SAS) hardware that enables the node hardware to use the drives of the expansion enclosure.

expansion enclosure
: A hardware unit that includes enclosure chassis, expansion canisters, drives, and function that allows extra drives to be connected.

export
: See share.

extended count key data (ECKD)
: An extension of the count-key-data (CKD) architecture. It includes additional commands that can be used to improve performance.

extended unique identifier (EUI)
: A unique iSCSI name that identifies an iSCSI target adapter or an iSCSI initiator adapter as defined by the iSCSI standard (RFC 3722).

extent
: A unit of data that manages the mapping of data between managed disks (MDisks) and volumes.

extent type
: See storage architecture type.

external storage
: Managed disks (MDisks) that are SCSI logical units presented by storage systems that are attached to and managed by the clustered system.

F

F_port
: See fabric port.

fabric port (F_port)
: An access point that is part of a Fibre Channel fabric. An F_port on a Fibre Channel fabric connects to a node's node port (N_port).

failback
: The restoration of an appliance to its initial configuration after detection and repair of a failed network or component.

failover
: An automatic operation that switches to a redundant or standby system or node in the event of a software, hardware, or network interruption. See also failback.

fan unit
: See power supply unit.

FB
: See fixed block.

FBA
: See fixed-block architecture.

FC
: See Fibre Channel.

FCIP
: See Fibre Channel over IP.

Fibre Channel (FC)
: A technology for transmitting data between computer devices. It is especially suited for attaching computer servers to shared storage devices and for interconnecting storage controllers and drives. See also zoning.

Fibre Channel extender
: A device used to extend a Fibre Channel link over a greater distance than is supported by the standard, usually a number of miles or kilometers. Devices must be deployed in pairs at each end of a link.

Fibre Channel over IP (FCIP)
: A network storage technology that combines the features of the Fibre Channel Protocol and the Internet Protocol (IP) to connect distributed SANs over large distances.

Fibre Channel port fan-in
: The number of hosts that can detect any one port.

field-replaceable unit (FRU)
: An assembly that is replaced in its entirety when any one of its components fails.

file module
: 1. A machine that provides file systems to network users. A file module must be provided with block volumes to store the data for these file systems. It also provides system management functions.
  2. A component that provides file systems to network users. A file module must be provided with storage for the file systems.

file set
: A subset of a file system that provides granularity of functions such as snapshots or quotas within the file system.

file system (FS)
: A collection of files and certain attributes associated with those files.

file system storage
: Data storage that is organized into files and directories.

File Transfer Protocol (FTP)
: In TCP/IP, an application layer protocol that uses TCP and Telnet services to transfer bulk-data files between machines or hosts.

fixed block (FB)
: See fixed-block architecture.

fixed-block architecture (FBA)
: An architecture for a virtual device that specifies the format of and access mechanisms for the virtual data units on the device. The virtual data unit is a block. All blocks on the device are the same size (fixed size). The system can access them independently. See also data record, storage architecture type.

fix procedure
: A maintenance procedure that runs within the product application and provides step-by-step guidance to resolve an error condition.

FlashCopy
: Pertaining to a point-in-time copy where a virtual copy of a volume is created. The target volume maintains the contents of the volume at the point in time when the copy was established. Any subsequent write operations to the source volume are not reflected on the target volume.

FlashCopy mapping
: A continuous space on a direct-access storage volume, which is occupied by or reserved for a particular data set, data space, or file.

FlashCopy relationship
: See FlashCopy mapping.

FlashCopy service
: A copy service that duplicates the contents of a source volume on a target volume. In the process, the original contents of the target volume are lost. See also point-in-time copy.

flash drive
: A data storage device that uses solid-state memory to store persistent data.

flash module
: A unit of the system that contains flash memory.

flush-through mode
: See write-through mode.

frame
: The hardware support structure, covers, and all electrical parts mounted therein that are packaged as one entity for shipping.

FRU
: See field-replaceable unit.

FS
: See file system.

FTP
: See File Transfer Protocol.

FTP server
: A software or hardware server that responds to File Transfer Protocol.

G

gateway
: An entity that operates above the link layer and converts, when required, the interface and protocol used by one network into those used by another distinct network.

GB
: See gigabyte.

GBIC
: See gigabit interface converter.

General Parallel File System (GPFS)
: A high-performance shared-disk file system that can provide data access from nodes in a clustered system environment.

gigabit interface converter (GBIC)
: An encoding/decoding device that is a class-1 laser component assembly with transmitting and receiving receptacles that connect to fiber-optic cables. GBICs perform a serial optical-to-electrical and electrical-to-optical conversion of the signal. The GBICs in the switch can be hot-swapped.

gigabyte (GB)
: For processor storage, real and virtual storage, and channel volume, two to the power of 30 or 1,073,741,824 bytes. For disk storage capacity and communications volume, 1,000,000,000 bytes.

Global Mirror
: 1. An asynchronous copy service that enables host data on a volume to be mirrored over long distances to a volume in a remote location.
  2. A method of an asynchronous replication that maintains data consistency across multiple volumes within or across multiple systems. Global Mirror is generally used where distances between the source site and target site cause increased latency beyond what the application can accept.

GPFS
: See General Parallel File System.

GPFS cluster
: A system of nodes that are defined as being available for use by GPFS file systems. See also cluster, cluster, clustered system.

GPFS snapshot
: A point-in-time copy of a file system or file set.

grain
: In a FlashCopy bitmap, the unit of data represented by a single bit.

graphical user interface (GUI)
: A computer interface that presents a visual metaphor of a real-world scene, often of a desktop, by combining high-resolution graphics, pointing devices, menu bars and other menus, overlapping windows, icons and the object-action relationship.

GUI
: See graphical user interface.

H

hardcoded
: Pertaining to software instructions that are statically encoded and not intended to be altered.

hard disk drive
: See drive.

Hardware Management Console (HMC)
: A system that controls managed systems, including the management of logical partitions and use of Capacity Upgrade on Demand. Using service applications, the HMC communicates with managed systems to detect and consolidate information, which is then sent to IBM for analysis.

HBA
: See host bus adapter.

heartbeat

A signal that one entity sends to another to convey that it is still active.

HMC
: See Hardware Management Console.

hop
: One segment of a transmission path between adjacent nodes in a routed network.

host
: A computer that is connected to a network and that provides an access point to that network. The host can be a client, a server, or both a client and server simultaneously. See also client.

host bus adapter (HBA)
: An interface card that connects a host bus, such as a peripheral component interconnect (PCI) bus, to the storage area network.

host group
: See host.

host ID
: A numeric identifier assigned to a group of host Fibre Channel ports for the purpose of logical unit number (LUN) mapping. For each host ID, there is a separate mapping of Small Computer System Interface (SCSI) IDs to volumes.

host interface card
: See interface card.

host mapping
: The process of controlling which hosts have access to specific volumes within a clustered system.

host zone
: A zone that is defined in the storage area network (SAN) fabric in which the hosts can address the system.

hot-spare
: Pertaining to redundant hardware (such as an adapter, a disk, a drive, or a server) that is installed and available in the event of a hardware failure.

HTTPS
: See Hypertext Transfer Protocol Secure.

hub
: A point, or piece of hardware, that connects multiple devices in a network.

HyperSwap®
: Pertaining to a function that provides continuous, transparent availability against storage errors and site failures, and is based on synchronous replication.

Hypertext Transfer Protocol Secure (HTTPS)
: An Internet protocol that is used by web servers and web browsers to transfer and display hypermedia documents securely across the Internet.

I

ID
: See identifier.

identifier (ID)
: A sequence of bits or characters that identifies a user, program, device, or system to another user, program, device, or system.

idling
: The status of a pair of volumes that have a defined copy relationship for which no copy activity has yet been started.

illegal configuration
: A configuration that will not operate and will generate an error code to indicate the cause of the problem.

image mode
: An access mode that establishes a one-to-one mapping of extents in the storage pool with the extents in the volume. See also access mode, managed mode, unconfigured mode.

image volume
: A volume in which there is a direct block-for-block translation from the managed disk (MDisk) to the volume.

IML
: See initial microcode load.

initial microcode load (IML)
: The loading of microcode from a storage medium into memory.

initiator
: The system component that originates an I/O command over an I/O bus or network. I/O adapters and network interface cards are typical initiators. See also target.

inode
: The internal structure that describes the individual files on AIX®, UNIX, or Linux® systems. An inode contains the node, type, owner, and location of a file.

input/output (I/O)
: Pertaining to a device, process, channel, or communication path involved in data input, data output, or both.

input/output operations per second (IOPS)
: A standard computing benchmark used to determine the best configuration settings for servers.

instance
: 1. In object-oriented programming, a region of storage that contains a value or group of values.
  2. A specific occurrence of an object that belongs to a class. See also object.

integrity
: In computer security, assurance that the information that arrives at a destination is the same as the information that was sent.

interface card
: An optional part of a node canister that provides the system with additional host and storage connectivity options.

interface node
: A node that connects a system to an Internet Protocol (IP) network for file-serving capabilities by using service protocols.

internal storage
: Array managed disks (MDisks) and drives that are held in enclosures and in nodes that are part of the clustered system.

Internet Protocol (IP)
: A protocol that routes data through a network or interconnected networks. This protocol acts as an intermediary between the higher protocol layers and the physical network.

Internet Storage Name Service Protocol (iSNS Protocol)
: A protocol that is used by a host system to manage iSCSI targets and iSCSI discovery. iSCSI initiators use the iSNS Protocol to locate the appropriate storage resources.

inter-switch link (ISL)
: The physical connection that carries a protocol for interconnecting multiple routers and switches in a storage area network (SAN).

I/O
: See input/output.

I/O group
: A collection of volumes and node relationships that present a common interface to host systems. Each pair of nodes is known as an input/output (I/O) group

IOPS
: See input/output operations per second.

I/O throttling rate
: The maximum rate at which an I/O transaction is accepted for a volume.

IP
: See Internet Protocol.

IP address
: A unique address for a device or logical unit on a network that uses the Internet Protocol standard.

IQN
: See iSCSI qualified name.

iSCSI alias
: An alternative name for the iSCSI-attached host.

iSCSI name
: A name that identifies an iSCSI target adapter or an iSCSI initiator adapter. An iSCSI name can be an iSCSI qualified name (IQN) or an extended-unique identifier (EUI). Typically, this identifier has the following format: iqn.datecode.reverse domain.

iSCSI qualified name (IQN)
: A unique name that identifies an iSCSI target adapter or an iSCSI initiator adapter as defined by the iSCSI standard (RFC 3722).

ISL
: See inter-switch link.

ISL hop
: Considering all pairs of node ports (N_ports) in a fabric and measuring distance only in terms of inter-switch links (ISLs) in the fabric, the number of ISLs traversed on the shortest route between the pair of nodes that are farthest apart in the fabric.

iSNS Protocol
: See Internet Storage Name Service Protocol.

J

JBOD
: See just a bunch of disks.

just a bunch of disks (JBOD)
: Hard disks that haven't been configured according to the RAID system to increase fault tolerance and improve data access performance.

L

latency
: The time interval between the initiation of a send operation by a source task and the completion of the matching receive operation by the target task. More generally, latency is the time between a task initiating data transfer and the time that transfer is recognized as complete at the data destination.

LBA
: See logical block address.

LDAP
: See Lightweight Directory Access Protocol.

least recently used (LRU)
: Pertaining to an algorithm used to identify and make available the cache space that contains the data that was least recently used.

Lightweight Directory Access Protocol (LDAP)
: An open protocol that uses TCP/IP to provide access to directories that support an X.500 model and that does not incur the resource requirements of the more complex X.500 Directory Access Protocol (DAP). For example, LDAP can be used to locate people, organizations, and other resources in an Internet or intranet directory.

line card
: See blade server.

local fabric
: Storage area network (SAN) components (such as switches and cables) that connect the components (nodes, hosts, switches) of the local clustered system.

local/remote fabric interconnect
: The storage area network (SAN) components that are used to connect the local and remote fabrics.

logical block address (LBA)
: The block number on a disk.

logical drive
: See volume.

logical unit (LU)
: An entity to which Small Computer System Interface (SCSI) commands are addressed, such as a volume or managed disk (MDisk).

logical unit number (LUN)
: In the Small Computer System Interface (SCSI) standard, a unique identifier used to differentiate devices, each of which is a logical unit (LU).

longitudinal redundancy check (LRC)
: A method of error-checking during data transfer that involves checking parity on a row of binary digits that are members of a set that form a matrix.

LRC
: See longitudinal redundancy check.

LRU

See least recently used.

LU
: See logical unit.

LUN
: See logical unit number.

LUN masking
: A process where a host object can detect more LUNs than it is intended to use, and the device-driver software masks the LUNs that are not to be used by this host.

M

MAC
: See Media Access Control.

managed disk (MDisk)
: A Small Computer System Interface (SCSI) logical unit (LU) that a Redundant Array of Independent Disks (RAID) controller provides and a clustered system manages. The MDisk is not visible to host systems on the storage area network (SAN).

managed disk group
: See storage pool.

managed mode
: An access mode that enables virtualization functions to be performed. See also access mode, image mode, unconfigured mode.

Managed Object Format (MOF)
: A language for defining Common Information Model (CIM) schemas.

management node
: A node that is used for configuring, administering, and monitoring a system.

master volume
: In most cases, the volume that contains a production copy of the data and that an application accesses. See also auxiliary volume, relationship.

MB
: See megabyte.

MBps
: See megabytes per second.

MDisk
: See managed disk.

Media Access Control (MAC)
: In networking, the lower of two sublayers of the Open Systems Interconnection model data link layer. The MAC sublayer handles access to shared media, such as whether token passing or contention will be used.

megabyte (MB)
: For processor storage, real and virtual storage, and channel volume, 2 to the 20th power or 1,048,576 bytes. For disk storage capacity and communications volume, 1,000,000 bytes.

megabytes per second (MBps)
: A unit of data transfer rate equal to 1024 * 1024 bytes.

mesh
: A network topology in which devices are connected with many redundant interconnections between network nodes. Every node has a connection to every other node in the network.

message
: An informational event that does not require user action. See also event.

method
: A way to implement a function on a class.

Metro Global Mirror
: A cascaded solution where Metro Mirror synchronously copies data to the target site. This Metro Mirror target is the source volume for Global Mirror that asynchronously copies data to a third site. This solution has the potential to provide a disaster recovery with no data loss at Global Mirror distances when the intermediate site does not participate in the disaster that occurs at the production site.

Metro Mirror
: 1. A function of the remote mirror and copy feature that constantly updates a secondary copy of a volume to match changes made to a source volume. See also capacity licensing, Global Mirror.
  2. A method of synchronous replication that maintains data consistency across multiple volumes within the system. Metro Mirror is generally used when the write latency caused by the distance between the source site and target site is acceptable to application performance.

mirrored volume
: A volume with two volume copies.

MOF
: See Managed Object Format.

N

N_port
: See node port.

namespace
: The scope within which a Common Information Model (CIM) schema applies.

NAS
: See network-attached storage.

nearline
: Pertaining to a type of storage in which data is available in a short amount of time, but not instantly.

nearline SAS drive
: A drive that combines the high capacity data storage technology of a Serial Advanced Technology Attachment (SATA) drive with the benefits of a serial-attached SCSI (SAS) interface for improved connectivity.

network-attached storage (NAS)
: A task-optimized storage device directly attached to a network that operates independently of the general-purpose file servers.

Network File System (NFS)
: A protocol that allows a computer to access files over a network as if they were on its local disks.

network interface controller (NIC)
: Hardware that provides the interface control between system main storage and external high-speed link (HSL) ports.

Network Shared Disk (NSD)
: A component for cluster-wide disk naming and access.

NFS
: See Network File System.

NIC
: See network interface controller.

node
: A single processing unit within a system. For redundancy, nodes are deployed in pairs to make up a system.

node canister
: A hardware unit that includes the node hardware, fabric and service interfaces, and serial-attached SCSI (SAS) expansion ports.

node port (N_port)
: A port that connects a node to a fabric or to another node. An N_port connects to a fabric port (F_port) or to the N_port of another node. An N_port handles creation, detection, and flow of message units to and from the connected systems. N_ports are end points in point-to-point links.

node rescue
: The process by which a node that has no valid software installed on its hard disk drive can copy software from another node connected to the same Fibre Channel fabric.

NSD
: See Network Shared Disk.

O

object
: In object-oriented design or programming, a concrete realization (instance) of a class that consists of data and the operations associated with that data. An object contains the instance data that is defined by the class, but the class owns the operations that are associated with the data.

object model
: An abstraction of a system's implementation.

object name
: An object that consists of a namespace path and a model path. The namespace path provides access to the Common Information Model (CIM) implementation managed by the CIM Agent, and the model path provides navigation within the implementation.

object path
: See object name.

operating set
: The set of nodes that are operating together to deliver storage services.

oversubscription
: The ratio of the sum of the traffic that is on the initiator network-device connections to the traffic that is on the most heavily loaded inter-switch links (ISLs), where more than one ISL is connected in parallel between these switches. The concept of oversubscription assumes a symmetrical network and a specific workload that is applied equally from all initiators and sent equally to all targets. See also symmetrical network.

P

partition
: A logical division of storage on a fixed disk. See also virtualization.

partner node
: The other node that is in the I/O group to which this node belongs.

partnership
: In Metro Mirror or Global Mirror operations, the relationship between two clustered systems. In a clustered-system partnership, one system is defined as the local system and the other system as the remote system.

pend
: A state between online and offline in which an object is making a transition from online to offline or from offline to online. The pend state is initiated when a pend event is sent by an online object or when an add event is sent by an offline object. See also un-pend.

point-in-time copy
: The instantaneous copy that the FlashCopy service makes of the source volume. See also FlashCopy service.

pool
: See storage pool.

pool pair
: Two storage pools that are required to balance workload. Each storage pool is controlled by a separate node.

port
: The physical entity within a host, system, or storage system that performs the data communication (transmitting and receiving) over the Fibre Channel.

POST
: See power-on self-test.

power-on self-test (POST)
: A series of internal diagnostic tests activated each time the system power is turned on.

power supply
: See power supply unit.

power supply unit (PSU)
: A component that supplies power to other components in a system.

primary volume
: In a stand-alone Metro Mirror or Global Mirror relationship, the target of write operations issued by the host application. See also relationship.

property

In the Common Information Model (CIM), an attribute that is used to characterize instances of a class.

PSU
: See power supply unit.

Q

qualifier

1. A value that provides additional information about a class, association, indication, method, method parameter, instance, property, or reference.
2. A modifier that makes a name unique.

queue depth
: The number of input/output (I/O) operations that can be run in parallel on a device.

quorum disk
: A disk that contains a reserved area that is used exclusively for system management. The quorum disk is accessed when it is necessary to determine which half of the clustered system continues to read and write data. Quorum disks can either be MDisks or drives.

quorum index
: The pointer that indicates the order used to resolve a tie. Nodes attempt to lock the first quorum disk (index 0), followed by the next disk (index 1), and finally the last disk (index 2). The tie is broken by the node that locks them first.

quota
: The amount of disk space and number of files and directories assigned as upper limits for a specified user, group of users, or file set.

R

rack
: A free-standing structure that can hold multiple servers, storage systems, chassis, switches, and other devices.

RAID
: See Redundant Array of Independent Disks.

RAID 0
: A data striping technique, which is commonly called RAID Level 0 or RAID 0 because of its similarity to common, RAID, data-mapping techniques. It includes no data protection, however, so, strictly speaking, the appellation RAID is a misnomer. RAID 0 is also known as data striping.

RAID 1
: A form of storage array in which two or more identical copies of data are maintained on separate media.

RAID 10
: A collection of two or more physical drives that present to the host an image of one or more drives. In the event of a physical device failure, the data can be read or regenerated from the other drives in the RAID due to data redundancy.

RAID 5
: A form of parity RAID in which the disks operate independently, the data stripe size is no smaller than the exported block size, and parity check data is distributed across the array's disks.

RAID 6
: A form of RAID that can continue to process read and write requests to all of an array's virtual disks in the presence of two concurrent disk failures.

RAID controller
: See node canister.

RAID level
: The level of protection provided by the specific techniques of striping, mirroring, or parity used by a Redundant Array of Independent Disks (RAID).

RAID type
: See RAID level.

real capacity
: The amount of storage that is allocated to a volume copy from a storage pool.

Recovery Guru
: See fix procedure.

redundant ac-power switch
: A device that provides input power redundancy by attaching a device to two independent power sources. If the main source becomes unavailable, the redundant ac-power switch automatically provides power from a secondary (backup) source. When power is restored, the redundant ac-power switch automatically changes back to the main power source.

Redundant Array of Independent Disks (RAID)
: A collection of two or more physical disk drives that present to the host an image of one or more logical disk drives. In the event of a physical device failure, the data can be read or regenerated from the other disk drives in the array due to data redundancy. See also array.

redundant SAN
: A storage area network (SAN) configuration in which any single component might fail, but connectivity between the devices within the SAN is maintained, possibly with degraded performance. This configuration is normally achieved by splitting the SAN into two independent, counterpart SANs. See also counterpart SAN.

rejected
: Pertaining to a status condition that describes a node that the clustered-system software has removed from the working set of nodes in the clustered system.

relationship
: In Metro Mirror or Global Mirror, the association between a master volume and an auxiliary volume. These volumes also have the attributes of a primary or secondary volume. See also auxiliary volume, master volume, primary volume, secondary volume.

remote copy

1. See Global Mirror.
2. See Metro Mirror.

remote fabric
: In Global Mirror, the storage area network (SAN) components (switches and cables) that connect the components (nodes, hosts and switches) of the remote clustered system.

root squash
: In the Network File System (NFS) Protocol, a reduction of the access rights for the remote superuser (root) when using identity authentication. The local user is the same as the remote user.

S

SAN
: See storage area network.

SAS
: See serial-attached SCSI.

schema
: A group of object classes defined for and applicable to a single namespace.

SCP
: See Secure Copy Protocol.

SCSI
: See Small Computer System Interface.

SCSI back-end layer
: The layer in a Small Computer System Interface (SCSI) network that performs the following functions: controls access to individual storage systems that are managed by the clustered system; receives requests from the virtualization layer, processes them, and sends them to managed disks; and addresses SCSI-3 commands to the storage systems on the storage area network (SAN).

SCSI front-end layer
: The layer in a Small Computer System Interface (SCSI) network that receives I/O commands from hosts and provides the SCSI-3 interface to hosts. SCSI logical unit numbers (LUNs) are mapped to volumes in this layer as well. Thus, the layer converts SCSI read and write commands that are addressed to LUNs into commands that are addressed to specific volumes.

secondary volume
: Pertinent to remote copy, the volume in a relationship that contains a copy of data written by the host application to the primary volume. See also relationship.

Secure Copy Protocol (SCP)
: The secure transfer of computer files between a local and a remote host or between two remote hosts, using the Secure Shell (SSH) protocol.

Secure Sockets Layer (SSL)
: A security protocol that provides communication privacy. With SSL, client/server applications can communicate in a way that is designed to prevent eavesdropping, tampering, and message forgery.

sequential volume
: A volume that uses extents from a single managed disk (MDisk).

serial-attached SCSI (SAS)
: A data-transfer technology that moves data to and from computer storage devices. Serial-attached SCSI uses a point-to-point serial protocol, which replaces the traditional, parallel SCSI bus technology.

Server Message Block (SMB)
: A protocol that manages requests and responses in a client/server environment so that clients on a network can share files, directories, and devices. See also Common Internet File System, Server Message Block 2.0.

Server Message Block 2.0
: A higher performing, more scalable version of Server Message Block (SMB). This protocol can send multiple commands in the same packet and uses larger buffer sizes. See also Common Internet File System, Server Message Block.

service assistant
: A user interface that services hardware independent of the storage system.

Service Location Protocol (SLP)
: An Internet protocol that identifies and uses network hosts without having to designate a specific network host name.

SFP transceiver
: See small-form-factor pluggable transceiver.

share
: A file system, file set, or directory that has been made accessible to authorized remote clients by using supported services.

Simple Network Management Protocol (SNMP)
: A set of protocols for monitoring systems and devices in complex networks. Information about managed devices is defined and stored in a Management Information Base (MIB).

SLP
: See Service Location Protocol.

Small Computer System Interface (SCSI)
: An ANSI-standard electronic interface that allows personal computers to communicate with peripheral hardware, such as disk drives, tape drives, CD-ROM drives, printers, and scanners faster and more flexibly than previous interfaces. See also target.

small-form-factor pluggable transceiver (SFP transceiver)
: An optical transceiver used to convert signals between optical fiber cables and switches.

SMB
: See Server Message Block.

SMI-S
: See Storage Management Initiative Specification.

snapshot
: An image backup type that consists of a point-in-time view of a volume.

SNIA
: See Storage Networking Industry Association.

SNMP
: See Simple Network Management Protocol.

solid-state drive
: 1. See flash drive.
  2. A data storage device that uses solid-state memory to store persistent data.

space efficient
: See thin provisioning.

space-efficient VDisk
: See thin-provisioned volume.

spare
: An extra storage component, such as a drive or tape, that is predesignated for use as a replacement for a failed component.

spare goal
: The optimal number of spares that are needed to protect the drives in the array from failures. The system logs a warning event when the number of spares that protect the array drops below this number.

SSD
: See solid-state drive.

SSL
: See Secure Sockets Layer.

stand-alone relationship
: In FlashCopy, Metro Mirror, and Global Mirror, relationships that do not belong to a consistency group and that have a null consistency-group attribute.

statesave
: Binary data collection that is used in problem determination.

storage architecture type (storage type)
: The type of storage architecture, either count key data (CKD) or fixed block (FB), for which an array, pool, or volume is provisioned. See also count key data, fixed-block architecture.

storage area network (SAN)
: A dedicated storage network tailored to a specific environment, combining servers, systems, storage products, networking products, software, and services.

storage controller enclosure
: See control enclosure.

storage enclosure
: An external storage device that is connected directly to a storage system. Storage enclosures provide storage capacity to the system in the form of hard disk drives (HDDs), flash drives or flash modules.

storage expansion enclosure
: See expansion enclosure.

Storage Management Initiative Specification (SMI-S)
: A design specification developed by the Storage Networking Industry Association (SNIA) that specifies a secure and reliable interface with which storage management systems (SMSs) can identify, classify, monitor, and control physical and logical resources in a storage area network (SAN). The interface integrates the various devices to be managed in a SAN and the tools used to manage them.

Storage Manager Enterprise Management window
: See graphical user interface.

Storage Networking Industry Association (SNIA)
: An alliance of computer vendors and universities that focus on developing and promoting industry standards for storage networks.

storage node
: A component of a storage system that provides internal storage or a connection to one or more external storage systems.

storage pod
: A subcomponent of a network-attached storage (NAS) system that consists of two or more storage nodes and one or more supported storage systems.

storage pool
: 1. A collection of storage capacity that provides the capacity requirements for a volume.
  2. A collection of storage that identifies an underlying set of resources. These resources provide the capacity and management requirements for a volume or set of volumes.

storage subsystem
: See system.

storage subsystem controller
: See node canister.

storage system
: 1. A system that provides persistent storage within a network. A storage system can include facilities for host attachment, user role authentication, a command-line interface (CLI), a graphical user interface (GUI), and storage devices that most often include Redundant Array of Independent Disks (RAID) controllers. It might also include agents for enabling third-party management software to monitor or manage the storage devices.
  2. A device, such as a Redundant Array of Independent Disks (RAID) controller, that creates and manages other storage devices.

storage type
: See storage architecture type.

strand
: The serial-attached SCSI (SAS) connectivity of a set of drives within multiple enclosures. The enclosures can be either control enclosures or expansion enclosures.

striped
: Pertaining to a volume that is created from multiple managed disks (MDisks) that are in the storage pool. Extents are allocated on the MDisks in the order specified.

subnet
: See subnetwork.

subnetwork (subnet)
: A network that is divided into smaller independent subgroups, which still are interconnected.

switch
: A network infrastructure component to which multiple nodes attach. Unlike a hub, a switch typically has internal bandwidth that is a multiple of link bandwidth and the ability to rapidly switch a node connection from one to another. A typical switch can accommodate several simultaneous full-link bandwidth transmissions between different pairs of nodes.

symmetrical network
: A network in which all the initiators are connected at the same level and all the controllers are connected at the same level. See also oversubscription.

symmetric virtualization

A virtualization technique in which the physical storage, in the form of a Redundant Array of Independent Disks (RAID), is split into smaller chunks of storage known as extents. These extents are then concatenated, using various policies, to make volumes. See also asymmetric virtualization.

synchronous replication

1. A type of replication in which the application write operation is made to both the source volume and target volume before control is given back to the application. See also asynchronous replication.
2. A type of replication that delivers updates continuously and within the scope of source transactions. See also asynchronous replication.

syslog
: A standard for transmitting and storing log messages from many sources to a centralized location to enhance system management.

system
: A functional unit, consisting of one or more computers and associated software, that uses common storage for all or part of a program and also for all or part of the data necessary for the execution of the program. A system can be a stand-alone unit, or it can consist of multiple connected units. See also clustered system.

system drive
: See boot drive.

T

target
: The program or system to which a request for files or processing is sent.

thin-provisioned volume
: A volume that allocates storage when data is written to it.

thin provisioning
: The ability to define a storage unit (full system, storage pool, volume) with a logical capacity size that is larger than the physical capacity assigned to that storage unit. See also volume.

throughput
: A measure of the amount of information transmitted over a network in a given period of time. Throughput is generally measured in bits per second (bps), kilobits per second (Kbps), or megabits per second (Mbps).

trigger
: To initiate or reinitiate copying between a pair of volumes that have a copy relationship.

U

UID
: See unique identifier.

unconfigured mode
: A mode in which I/O operations cannot be performed. See also access mode, image mode, managed mode.

uninterruptible power supply (UPS)
: A source of power from a battery installed between the commercial power and the system that keeps the system running, if a commercial power failure occurs, until it can complete an orderly end to system processing.

unique identifier (UID)
: An identifier that is assigned to storage-system logical units when they are created. It is used to identify the logical unit regardless of the logical unit number (LUN), the status of the logical unit, or whether alternate paths exist to the same device. Typically, a UID is used only once.

un-pend
: An event sent by a object in the pend state that results in the moving of the object to the online state. See also pend.

UPS
: See uninterruptible power supply.

USB key
: See flash drive.

user role
: An identifier that is assigned to a user that defines the set of tasks that the user can perform.

V

VDisk
: See volume.

VDisk-to-host mapping
: See host mapping.

VDP
: See vital product data.

virtual capacity
: The amount of storage that is available. In a thin-provisioned volume, the virtual capacity can be different from the real capacity. In a standard volume, the virtual capacity and real capacity are the same.

virtual disk
: See volume.

virtualization
: In the storage industry, a concept in which a pool of storage is created that contains several storage systems. Storage systems from various vendors can be used. The pool can be split into volumes that are visible to the host systems that use them. See also capacity licensing.

virtualized storage
: Physical storage that has virtualization techniques applied to it by a virtualization engine.

virtual storage area network (VSAN)
: A fabric within the storage area network (SAN).

vital product data (VDP, VPD)
: Information that uniquely defines system, hardware, software, and microcode elements of a processing system.

volume

1. The representation of an Small Computer System Interface (SCSI) disk that is presented to a host system. See also thin provisioning.
2. A fixed amount of physical or virtual storage on a data storage medium.

volume access set
: The set of I/O groups that allows host access to a volume. This set can optionally include the caching I/O group.

volume copy
: A physical copy of the data that is stored on a volume. Mirrored volumes have two such copies. Nonmirrored volumes have one copy.

VPD
: See vital product data.

VSAN
: See virtual storage area network.

W

worldwide name (WWN)
: A 64-bit, unsigned name identifier that is unique.

worldwide node name (WWNN)
: A unique 64-bit identifier for a host containing a Fibre Channel port. See also worldwide port name.

worldwide port name (WWPN)
: A unique 64-bit identifier associated with a Fibre Channel adapter port. The WWPN is assigned in an implementation-independent and protocol-independent manner. See also worldwide node name.

write-through mode
: A process in which data is written to a storage device at the same time as the data is cached.

WWN
: See worldwide name.

WWNN
: See worldwide node name.

WWPN
: See worldwide port name.

Z

z Global Mirror
: A method of an asynchronous replication function that maintains data consistency across multiple volumes that are attached to a z/OS® system. Time-based data consistency is maintained through the Data Facility Storage Management Subsystem (DFSMS) system data mover (SDM) component.

zone
: A collection of Fibre Channel device ports that are permitted to communicate with each other using the fabric. Types of device ports for zone purposes are node ports (N_ports) or node loop ports (NL_ports). Any two N_ports or NL_ports that are not members of at least one common zone are not permitted to communicate using the fabric. See also zoning.

zoning
: In Fibre Channel environments, the grouping of multiple ports to form a virtual, private, storage network. Ports that are members of a zone can communicate with each other, but are isolated from ports in other zones. See also Fibre Channel, zone.