

XIV Storage System



Tables of Contents

Welcome	1
Cross-system feature use in the IBM Spectrum Accelerate family	1
GDPR and IBM systems	4
IBM XIV Storage System Gen3 (11.x) documentation	5
Agreements and license information	5
Notices	5
Trademarks	7
Homologation statement	7
Safety and environmental notices	7
Safety notices and labels	7
Special caution and safety notices	8
Laser safety	8
Ladder usage	8
Fire suppression systems	8
Power cables	9
Leakage current	9
Site preparation	9
Environmental notices	10
Electromagnetic compatibility notices	10
Canada Notice	10
European Community and Morocco Notice	10
Germany Notice	10
Japan Electronics and Information Technology Industries Association (JEITA) Notice	11
Japan Voluntary Control Council for Interference (VCCI) Notice	12
Korea Notice	12
People's Republic of China Notice	12
Russia Notice	12
Taiwan Notice	12
United Kingdom Notice	13
United States Federal Communications Commission (FCC) Notice	13
Warranty information	13
What's new	13
Release notes	16
Release notes for IBM XIV Storage System	16
What's new in version 11.6.2.d	16
Hot upgrade from earlier versions	16
Management and solutions	17
Change log	17
11.6.2.d (August 2019)	18
11.6.2.c (August 2018)	18
11.6.2.b (May 2018)	19
11.6.2.a (July 2017)	19
11.6.2 (December 2016)	20
11.6.1.a (March 2016)	22
11.6.1 (December 2015)	22
11.6.0.b (June 2015)	23
Fixes and features of earlier versions	24
Limitations	24
Known issues	25
Release notes for IBM Hyper-Scale Manager	28
Release notes for IBM XCLI Utility	28
Overview	28
What's new in 5.5.4.1	28
Compatibility	28
Change log	29
5.5.4.1 (June 2022)	29
5.5.4 (May 2022)	29
5.2.0.a (July 2017)	29
5.2.0 (April 2017)	29
5.1.0 (November 2016)	29
Limitations	30
Known issues	30
Release notes for IBM XIV Management Tools	30
What's new in version 4.8.0.6	30
Compatibility	30
Requirements	30

XIV GUI requirements	31
IBM Hyper-Scale Manager requirements	32
Change log	32
Version 4.8.0.6 (July 2017)	33
Version 4.8.0.4 (December 2016)	33
Version 4.8.0.3 (July 2016)	33
Version 4.8.0.2 (April 2016)	33
Version 4.8.0.1 (January 2016)	34
Version 4.8 (December 2015)	34
Fixes and features of earlier versions	34
Limitations	34
Known issues	34
Release notes for IBM Remote Support Proxy	37
What's new in 1.3.2.2	37
Supported operating systems	38
Supported storage systems	38
Limitations	38
Known issues	38
Release notes PDF publications	38
Product overview	40
Introduction	41
Features and functionality	42
Hardware overview	42
Hardware components	43
Hardware enhancements	43
Management options	44
Reliability	44
Redundant components	44
Data mirroring	44
Self-healing mechanisms	44
Protected cache	45
Redundant power	45
SSD cache drives	45
Performance	45
Functionality	46
Snapshot management	46
Consistency groups for snapshots	46
Storage pools	46
Remote monitoring and diagnostics	46
SNMP	47
Multipathing	47
Automatic event notifications	47
Management through GUI and CLI	47
External replication mechanisms	47
Upgradability	47
Volumes and snapshots	47
Volume function and lifecycle	48
Support for Symantec Storage Foundation Thin Reclamation	48
Snapshot function and lifecycle	48
Creating a snapshot	49
Locking and unlocking snapshots	49
Duplicating a snapshot	50
Creating a snapshot of a snapshot	50
Formatting a snapshot or a snapshot group	50
Additional snapshot attributes	51
Redirect-on-Write (ROW)	51
Full Volume Copy	53
Restoring volumes and snapshots	53
Storage pools	54
Protecting snapshots at a storage pool level	55
Thin provisioning	55
Instant space reclamation	56
Consistency groups	56
Snapshot of a consistency group	57
Consistency group snapshot lifecycle	58
QoS performance classes	59
Connectivity with hosts	60
IP and Ethernet connectivity	60
Ethernet ports	61
IPv6 certification	61
Management connectivity	61
Field technician ports	61

Configuration guidelines summary	62
Host system attachment	62
Balanced traffic without a single point of failure	62
Dynamic rate adaptation	63
Attaching volumes to hosts	63
Excluding LUN0	63
Advanced host attachment	63
CHAP authentication of iSCSI hosts	63
Clustering hosts into LUN maps	64
Volume mapping exceptions	65
Supporting VMware extended operations	66
Writing zeroes	66
Hardware-assisted locking	66
Fast copy	66
IBM Real-time Compression with XIV	67
Planning for compression	68
Understanding compression rates, ratios and savings	68
Prerequisites and limitations	68
Estimating compression savings	69
Effective capacity	69
General compression saving guidelines	69
Estimating compression savings using XIV GUI	70
Estimating compression savings using IBM Comprestimator utility	70
Synchronous remote mirroring	71
Remote mirroring basic concepts	71
Remote mirroring operation	72
Configuration options	72
Volume configuration	72
Mixed configuration	73
Communication errors	73
Coupling activation	73
Synchronous mirroring statuses	73
Link status	74
Operational status	74
Synchronization status	74
Synchronization progress status	75
Last secondary timestamp	75
I/O operations	75
Synchronization process	75
State diagram	76
Coupling recovery	76
Uncommitted data	77
Constraints and limitations	77
Last-consistent snapshots	77
Secondary locked error status	78
Role switchover	78
Role switchover when remote mirroring is operational	78
Role switchover when remote mirroring is nonoperational	78
Switch secondary to primary	79
Secondary consistency	79
Switch primary to secondary	79
Resumption of remote mirroring after role change	80
Reconnection when both sides have the same role	80
Remote mirroring	80
Remote mirroring and consistency groups	81
Using remote mirroring for media error recovery	81
Supported configurations	81
I/O performance versus synchronization speed optimization	81
Implications regarding volume and snapshot management	81
Asynchronous remote mirroring	81
Features	82
Asynchronous remote mirroring terminology	83
Specifications	83
Technological overview	84
Replication scheme	84
Snapshot-based technology	85
Mirroring-special snapshots	85
Initializing the mirroring	86
Off-line replicating the master onto the slave	86
The sync job	86
Mirroring schedules and intervals	87

The mirror snapshot (ad-hoc sync job)	88
Determining replication and mirror states	88
RPO and RTO	88
Mirror status values	88
Definitions of mirror state and status	89
Determining mirror status	89
Mirrored consistency groups	91
Storage space required for the mirroring	91
Pool space depletion	91
Protecting snapshots using deletion priority	91
Pool space depletion on the master	93
Pool space depletion on the slave	94
Asynchronous mirroring process walkthrough	94
Step 1	95
Step 2	95
Step 3	95
Step 4	96
Step 5	96
Step 6	96
Step 7	97
Step 8	97
Step 9	98
Step 10	98
Step 11	98
Peers roles	99
Activating the mirroring	99
Link states	99
Activation states	99
Deactivating the coupling	100
Mirroring consistency groups	100
Setting a consistency group to be mirrored	101
Creating a mirrored consistency group	101
Adding a mirrored volume to a mirrored consistency group	102
Removing a volume from a mirrored consistency group	102
Multi-site mirroring	102
Multi-site mirroring terminology	102
Multi-site mirroring technological overview	103
IBM Hyper-Scale Mobility	104
The IBM Hyper-Scale Mobility process	104
Data-at-rest encryption	105
HIPAA compatibility	106
Management and monitoring	106
Event notification destinations	106
Event information	107
Event notification rules	107
Event notification gateways	108
User roles and permissions	108
User groups	109
Predefined users	109
User information	109
User authentication and access control	109
Native authentication	110
LDAP authentication	110
LDAP authentication logic	110
Multi-Tenancy	111
Working with multi-tenancy	113
Integration with ISV environments	113
VMware Virtual Volumes and XIV	114
IBM XIV support of Microsoft Azure Site Recovery (ASR)	114
Software upgrade	115
Preparing for upgrade	116
Remote support and proactive support	116
Learning and tutorials	117
Reinventing your storage reality with IBM XIV Storage System e-Learning course	117
Remote support e-Learning course	118
Planning	119
System specifications	120
Planning for compression	121
Understanding compression rates, ratios, and savings	121
Prerequisites and limitations	121
Estimating compression savings	122

Effective capacity	122
General compression saving guidelines	122
Estimating compression savings using XIV GUI	123
Estimating compression savings using IBM Comprestimator utility	123
Physical configuration and features	124
Rack components	124
Rack configuration features	125
Capacity	127
Processors and memory	128
Connectivity	128
Data and interface modules	129
Solid-state drive caching	130
Radio frequency identification device	130
Ruggedized rack	130
Rear-door heat exchanger	131
Weight-and-height-reduced shipping	132
Delivery and relocation requirements	132
Planning to receive delivery	132
Shipment weights and dimensions	133
Installation site requirements	133
Floor and space requirements	133
Raised or non-raised floor considerations	134
Floor-load requirements	135
Rack dimensions and service clearance requirements	135
Preparing for raised-floor installation and cabling	136
Preparing for non-raised-floor installation and cabling	136
Preparing for the ruggedized rack	136
Preparing for the rear-door heat exchanger	137
Power requirements	137
Power sources	138
Power consumption	138
Input voltages and frequencies	139
Main-power cables specifications	139
Emergency power-off considerations	140
Environmental requirements	141
Operating environment requirements	141
System environmental acclimation	141
Air circulation and cooling	142
Temperature thresholds and events	142
Contamination information	143
Acoustic declaration	143
Operating vibration requirements	144
Network and host connectivity requirements	144
Network connections	145
Management ports requirements	146
Network configurations	147
Fibre Channel network configurations	147
Ethernet (iSCSI) network configurations	149
Network cable requirements	150
Management-tool requirements	150
Requirements for IBM XIV Storage Management GUI, XCLI, and XIVTop	150
GUI specifications	150
Remote support and call home	151
Remote support connections	152
Remote access through a dedicated network connection	152
Remote access by modem	153
Call home configuration	153
Planning for the rear-door heat exchanger	154
Operating specifications	154
Performance	155
Preparing your site to provide water to the heat exchanger	156
Water specifications for the secondary cooling loop	156
Water delivery specifications for secondary loops	157
Raised floor hose requirements and management	161
Non-raised floor hose requirements and management	162
Secondary cooling loop parts and services information	163
Maintenance schedule	165
Administering users	165
User roles and permissions	108
User groups	109
Predefined users	109
User information	109

User authentication and access control	109
Native authentication	110
LDAP authentication	110
LDAP authentication logic	110
Managing	169
Managing systems using IBM Hyper-Scale Manager 5.x	169
Managing systems using IBM XIV Management Tools	169
Managing systems using the IBM XIV command-line interface (XCLI)	169
Operating	169
Powering on and off the XIV system	169
Powering on the XIV system, models 281x-114, 281x-116, 281x-214, and 281x-314, after shutdown	170
Verifying the system state	171
Powering off XIV system models 281x-114, 281x-116, 281x-214, and 281x-314	172
Installing and upgrading software	172
Non-disruptive code load	172
Preparing for code load	173
Remote support	174
Remote support and proactive support	174
Remote Support Proxy	174
Introducing the IBM Remote Support Proxy utility	174
Compatibility and requirements	175
Installation	175
Installing and configuring for the first time	175
Installing the software package	175
Configuring the Remote Support Proxy service	176
Starting the Remote Support Proxy service	178
Configuring storage systems	179
Upgrading and configuring an existing installation	179
Uninstalling the Remote Support Proxy	179
Management	180
Starting and stopping the Remote Support Proxy service	180
Monitoring status	180
Updating remote-support server information	180
Reference	181
Command-line reference	181
Overview of the command-line interface (CLI)	182
Overview of the XCLI utility	182
Installing and starting the XCLI	182
Starting the XCLI on a Windows system	183
Starting the XCLI on a Linux system	183
Exiting an interactive XCLI session	184
Using the CLI	184
Interactive and basic modes	184
Interactive mode features	185
Understanding the command syntax	185
Command syntax quick reference	186
Syntax example	186
Using identification parameters	186
Setting user and password parameters	187
Identifying and configuring a storage system	187
Configuration parameters	187
Creating a configuration file	188
Certificate management	188
Display options	189
Using display options	189
Table display options	190
Running commands in batch mode	191
Displaying CLI help and the XCLI utility version	192
Host and cluster management commands	192
Adding a host to a cluster	193
Creating a cluster	194
Deleting clusters	195
Listing clusters	195
Removing a host from a cluster	196
Renaming clusters	196
Adding a port to a host	197
Defining a new host	198
Deleting a host	199
Listing hosts	199

Listing ports	200
Removing a port from a host	201
Renaming a host	202
Updating a host definition	202
Mapping a volume to a host or cluster	203
Listing the mapping of volumes to hosts or clusters	205
Setting the special type of hosts or clusters	206
Listing hosts/clusters to which a volume is mapped	206
Unmapping a volume from a host or cluster	207
Setting the default idle time before unmapping a volume	208
Retrieving the default idle time before unmapping a volume	208
Creating a performance class	209
Deleting a performance class	209
Renaming a performance class	210
Listing details of performance classes	211
Adding a host to a performance class	211
Removing a host from its performance class	212
Adding a pool to a performance class	212
Removing a pool from its performance class	213
Adding a domain to a performance class	213
Removing a domain from its performance class	214
Setting the rate for a performance class	215
Listing host profiles	215
Updating the host profile	216
Removing the profile of the specified host	216
Enabling the host profiler	217
Disabling the host profiler	217
Creating an ALU LUN	217
Deleting an ALU LUN	218
Listing ALUs	219
Binding volumes to an ALU	219
Unbinding volumes from an ALU	220
Unbinding all volumes from an ALU	221
Listing all volumes bound to an ALU	221
Listing all ALUs and hosts to which the volume is bound	222
Volume management commands	222
Clearing reservations of a volume	223
Listing reservation keys	224
Listing volume reservations	224
Finding a volume by its internal index	225
Copying volumes	226
Copying volume differences	228
Creating a volume	229
Creating multiple volumes	231
Deleting a volume	232
Formatting a volume	233
Listing volumes	234
Locking a volume	235
Renaming a volume	236
Resizing a volume	237
Unlocking a volume	240
Compressing volumes	240
Managing volume compression and decompression	242
Deleting a compression or decompression request	243
Listing the volume compression activity	243
Setting a volume's external identifier	244
Clearing a volume's external identifier	245
Displaying the allocated partitions of a volume	245
Displaying the difference of allocated partitions between two volumes	246
Triggering compression estimation for a system volume	247
Removing a volume from the compression estimation queue	247
Displaying the volume's estimated compression ratio	248
Volume snapshot management commands	249
Changing a snapshot deletion priority	249
Creating a snapshot	250
Creating multiple snapshots	252
Deleting a snapshot	253
Duplicating a snapshot	253
Formatting a snapshot	255

Listing snapshot information	255
Restoring a volume from a snapshot	257
Consistency group management commands	258
Adding a volume to a consistency group	259
Creating consistency groups	261
Deleting a consistency group	262
Listing consistency groups	263
Removing a volume from a consistency group	263
Renaming a consistency group	265
Suspending I/O execution on a consistency group	265
Resuming I/O execution	266
Listing the status of consistency groups with paused I/O	267
Creating a cross-system consistency group	267
Associating an existing consistency group with a cross-system consistency group definition	268
Removing a consistency group from a cross-system consistency group	269
Adding a remote system name to a cross-system consistency group definition	270
Removing a remote system from a cross-system consistency group	270
Listing cross-system consistency group definitions	271
Retrieving remote systems in a specified cross-system consistency group	271
Deleting a cross-system consistency group	272
Listing cross-system consistency group definitions	273
Grouped pool management commands	273
Creating grouped pools	273
Listing grouped pools	274
Deleting a grouped pool	275
Renaming grouped pools	276
Moving capacity in a grouped pool	276
Snapshot set management commands	277
Snapshotting a consistency group	278
Changing a snapshot group deletion priority	280
Deleting a snapshot group	280
Disbanding a snapshot group	281
Duplicating a snapshot group	282
Formatting a snapshot group	283
Listing snapshot groups	284
Locking a snapshot group	284
Renaming a snapshot group	285
Restoring a consistency group from a snapshot group	286
Unlocking a snapshot group	287
Setting a snapshot group descriptor	287
Returning a snapshot group's descriptor	288
Storage pool management commands	289
Moving a consistency group between storage pools	289
Changing the pool limitation, performance class, or threshold parameters	290
Changing pool settings for snapshots	292
Creating storage pools	293
Deleting a storage pool	294
Listing storage pools	295
Listing thin pool compression savings	296
Renaming a storage pool	297
Resizing a storage pool	297
Moving a volume between storage pools	299
System management commands	300
Displaying current consumed capacity of the system	301
Displaying the values of configuration parameters	302
Setting configuration parameters	303
Testing the DNS	304
Displaying help	305
Displaying the current maintenance urgency	305
Shutting down the system	306
Listing the operational state	307
Displaying the system power management data	307
Displaying information about effective and physical capacity	308
Setting the system's soft capacity	308
Displaying the current time	309
Setting the system's time	309
Listing optional time zones	310
Setting the time zone	310
Printing the current system version	311

Displaying the values of VPD parameters	311
Setting VPD parameters	312
Displaying values of maintenance module parameters	313
Displaying the system's MIB file	314
Retrieving the electronic license acceptance status	315
Retrieving a fragment of the electronic license file	315
Accepting the electronic license agreement	316
Enabling command auditing	316
Disabling command auditing	317
Displaying the command audit state	317
Configuring audit servers	318
Checking the command audit state	318
Enabling CIM service	319
Disabling the CIM service	319
Displaying the CIM service state	320
Remote target connectivity commands	320
Setting the threshold of a link disruption duration that triggers an event	321
Updating the target's mirroring configuration	321
Activating connectivity to a remote target	322
Deactivating connectivity to a remote target	322
Defining connectivity to a remote target	323
Deleting connectivity to a remote target	324
Listing target connectivity definitions	325
Defining a remote target	326
Deleting a remote target	327
Listing remote targets	327
Allowing remote mirroring access	328
Activating a port	329
Adding a new port to a remote target	329
Deactivating a port	330
Deleting a port from a remote system	331
Listing the ports of a remote target	331
Renaming a remote target	332
Updating the target configuration	332
Remote mirroring commands	333
Canceling a snapshot mirror (ad hoc sync job)	334
Creating a snapshot mirror (ad hoc sync job)	335
Activating mirroring	338
Changing the RPO for local or remote system	339
Changing the designation of mirroring peers	340
Changing the mirroring schedule for remote slave peers	342
Changing the roles of a mirrored volume	343
Changing a mirroring schedule for local peers	345
Creating a mirroring definition	346
Deactivating mirroring	350
Deleting a remote mirroring definition	352
Viewing the mirroring status	353
Obtaining statistics on past sync jobs	355
Switching roles between master and slave	356
Retrieving RPO thresholds	358
Setting an RPO threshold	358
Changing the interval of a schedule	359
Creating a schedule object	360
Triggering a schedule	361
Deleting a schedule object	361
Listing a schedule object	362
Renaming a schedule	363
Viewing sync job status	363
Data migration commands	364
Activating data migration	364
Deactivating data migration	365
Defining data migration configuration	365
Deleting the data migration process	367
Listing data migration statuses	368
Testing the data migration definition	368
IBM Hyper-Scale Mobility commands	369
Creating an IBM Hyper-Scale Mobility relation	369
Activating a volume migration	373
Deactivating IBM Hyper-Scale Mobility migration	374

Aborting a defined or activated IBM Hyper-Scale Mobility process	375
Moving the IBM Hyper-Scale Mobility source volume to a Proxy state	375
Deleting an IBM Hyper-Scale Mobility relation	376
Listing the IBM Hyper-Scale Mobility status	377
Event handling commands	378
Generating a custom event	379
Generating a CSS product event	380
Defining a new event notification destination	380
Deleting a destination	382
Listing event notification destinations	383
Renaming a destination	384
Testing a destination	385
Updating an event notification destination	386
Adding a destination to a destination group	388
Creating a destination group	389
Updating an event notification destination group	389
Deleting a destination group	390
Listing destination groups	391
Removing a destination from a destination group	391
Renaming a destination group	392
Clearing alerting events	392
Listing events	393
Listing uncleared alerting events	395
Setting the threshold for event notification	395
Listing thresholds	396
Generating a custom event	397
Retrieving a maintenance module's heartbeat	397
Activating a rule	398
Creating event notification rules	398
Deactivating a rule	400
Deleting event notification rules	401
Listing event notification rules	402
Renaming event notification rules	402
Updating an event notification rule	403
Defining an SMS gateway	405
Deleting an SMS gateway	406
Listing SMS gateways	407
Prioritizing SMS gateways	407
Renaming an SMS gateway	408
Updating an SMS gateway	409
Defining a new SMTP gateway	409
Deleting an SMTP gateway	410
Listing SMTP gateways	411
Prioritizing SMTP gateways	412
Renaming an SMTP gateway	412
Updating an SMTP gateway	413
Generating an XMPNS user control event	414
Generating an XMPNS admin control event	414
IP configuration commands	415
Adding Ethernet ports to IP interfaces	416
Creating a new IP interface	416
Deleting IP interfaces	418
Listing IP interface configuration	418
Listing IP interface addresses	419
Showing the status and configuration of Ethernet ports	420
Removing Ethernet ports from IP interfaces	421
Renaming an IP interface	422
Printing the ARP database of an IP interface	422
Testing the traceroute to a remote IP	423
Testing the traceroute to a remote IP	423
Updating an IP interface	424
Defining a new IPSec connection	425
Updating an existing IPSec connection	426
Removing an existing IPSec connection	427
Listing IPSec connections	428
Listing IPSec tunnels	428
Connecting to a support center	429
Defining a support center	430
Deleting a support center	430

Disconnecting from a support center	431
Listing support centers	431
Listing the status of all support centers	432
Defining the automatic connection to a support center	432
Listing the configuration of automatically connected support centers	433
Creating a new IP access group	433
Removing an address from an IP access group	434
Adding a new address to an IP access group	435
Deleting an existing IP access group	435
Renaming an existing IP access group	436
Listing IP access groups	436
Listing IP access groups	437
PKI configuration commands	437
Listing PKI items	438
Generating a certificate signing request	438
Generating a private key and CSR	439
Deleting the PKI content	440
Changing a PKI symbolic name	440
Importing a signed certificate	441
Importing a PKCS#12 certificate	442
Displaying the details of a signed certificate	443
Updating a PKI certificate or services	444
Compression commands	445
Activating compression	445
Deactivating compression	446
Displaying the state of the estimated compression savings feature	446
Controlling the continuous display of the estimated compression savings	447
InfiniBand commands	447
Listing the configured InfiniBand ports	447
Listing the configured InfiniBand switches	448
Access control commands	450
Adding an access control definition	451
Deleting an access control definition	452
Listing access control definitions	452
Adding an LDAP server definition	453
Testing an LDAP configuration	454
Listing LDAP configuration parameters	455
Configuring LDAP in the system	456
Listing LDAP servers defined in the system	458
Listing LDAP server users	459
Listing LDAP-based authentication mode	459
Enabling or disabling LDAP-based authentication mode	460
Updating an LDAP server definition	461
Removing an LDAP server definition	462
Launching the ldapsearch utility	462
Defining a new user	464
Deleting a user	465
Adding users to user groups	465
Creating user groups	466
Deleting a user group	467
Listing user groups	468
Removing a user from a user group	468
Renaming user groups	469
Updating a user group	470
Listing users	471
Renaming users	471
Updating a user definition	472
Creating a new domain	473
Updating a domain definition	475
Renaming a domain	476
Deleting a domain	477
Listing domains	477
Listing users per domain	479
Listing objects in domains	479
Listing the global domain	480
Attaching an object to a domain	481
Disassociating object from a domain	482
Associating users to a domain	483
Removing a user from a domain	484

Adding a pool to a domain	485
Removing a pool from a domain	486
Moving a pool from one domain to another	487
Moving a grouped pool from one domain to another	488
Setting the domain attribute	490
Setting domain-related policies	490
Displaying domain-related policies	491
Specifying a user associated with IBM Hyper-Scale Manager	492
Retrieving the user associated with the IBM Hyper-Scale Manager	492
Setting the application administrator's scope of commands	493
Getting the application administrator's scope of commands	493
Fibre channel and iSCSI configuration and status commands	494
Discovering FC hosts	494
Changing FC port configuration	494
Listing FC ports	495
Listing connectivity to hosts	497
Hardware maintenance commands	497
Displaying an ATS configuration	498
Displaying the system's coupling facilities	500
Listing system components	501
Listing system components that require service	501
Listing InfiniBand host card adapters in the storage system	502
Listing CNA adapters in the system	503
Listing DIMMs in the system	504
Listing CPUs in the system	505
Listing maintenance modules in the system	506
Listing NICs in the system	507
Listing the internal temperature of modules	508
Monitoring the rebuilding or redistribution processes	510
Listing disk status	510
Listing module configuration	511
Checking serial connections of modules	512
Listing serial console statuses	513
Displaying the status of a UPS component	513
Listing service statuses	515
Listing PSUs in the system	516
Creating a trace snapshot	516
Listing trace snapshot on a module	517
Enabling XIV Support access	517
Disabling XIV Support access	518
Displaying the XIV Support window	519
Listing fans in the system	519
Listing SSDs used as flash cache	520
Getting the default state of the SSD caching	521
Setting the default for SSD caching	521
Overriding the SSD caching state	522
Displaying the system's average power consumption	523
Displaying the system's average temperature	523
Statistics commands	523
Getting performance statistics	523
Retrieving usage history	526
Metadata commands	527
Setting metadata	527
Deleting metadata	528
Listing metadata	528
Metadata-service commands	529
Creating a metadata service database	529
Deleting a metadata service database	530
Listing all metadata service databases	530
Adding a key to the metadata service	531
Deleting a key from the metadata service	531
Listing all keys in a database	532
Retrieving a metadata service key	532
Updating the metadata service key	533
Enabling the metadata service	533
Encryption enablement and support commands	534
Disabling encryption	534
Enabling encryption	535
Defining a keyserver	537

Removing a keyserver	538
Displaying keyserver status	539
Obtaining a new master key	539
Renaming a keyserver	541
Changing keyserver properties	541
Entering a recovery key	542
Generating recovery keys	544
Retrieving the security administrator's recovery key	545
Rekeying the security administrators	546
Displaying recovery key status	547
Recovering key verification	548
Recovering key share information	549
Finishing the recovery process	550
xmirror commands	550
Viewing xmirror status	551
Defining an xmirror relation	551
Activating an xmirror	553
Deactivating an xmirror	555
Deleting an xmirror relation	555
Renaming an xmirror relation	556
Changing roles in an xmirror relation	556
Defining a standby mirror in an xmirror relation	558
Security configuration commands	559
Listing configuration parameters for a communication protocol	559
Setting configuration parameters for a communication protocol	559
Events	560
Return codes	667
REST API reference	667
XIV Open API	667
XIV Open API overview	668
CIM agent overview	668
CIM agent components	669
CIM agent concepts	669
CIM agent security	670
CIM agent for theXIV system	670
CIM agent installation overview	670
Configuring the CIM agent	670
Enabling the CIM agent	671
XIV Open API component definitions	671
CIM agent communication concepts	672
CIM agent communication methods	672
GetClass	673
GetInstance	674
DeleteClass	674
DeleteInstance	675
CreateClass	675
CreateInstance	675
ModifyClass	675
ModifyInstance	676
EnumerateClasses	676
EnumerateClassNames	676
EnumerateInstances	677
EnumerateInstanceNames	677
ExecuteQuery	678
Associators	678
AssociatorNames	678
References	679
ReferenceNames	679
GetProperty	680
SetProperty	680
GetQualifier	680
SetQualifier	681
DeleteQualifier	681
EnumerateQualifiers	681
Return error codes	682
Functional diagrams of the CIM agent	683
Block Server performance profile	683
Block Server Performance object model	684
IBMTSDS_PerformanceStatisticsService. GetStatisticsCollection	685
Getting volume or host statistics with GetStatisticsCollection	685

Block Services profile	686
Block Services object model	686
IBMTSDS_StorageConfigurationService. CreateOrModifyStoragePool	689
IBMTSDS_StorageConfigurationService. DeleteStoragePool	689
IBMTSDS_StorageConfigurationService. CreateOrModifyElementsFromStoragePool	689
IBMTSDS_StorageConfigurationService. CreateOrModifyElementFromStoragePool	690
IBMTSDS_StorageConfigurationService. ReturnToStoragePool	690
IBMTSDS_StorageConfigurationService. ReturnElementsToStoragePool	690
IBMTSDS_PrimordialStoragePool. GetSupportedSizeRange	691
IBMTSDS_VirtualPool. GetSupportedSizeRange	691
Creating a ConcreteStoragePool	691
Creating a StorageVolume	691
iSCSI Target Ports profile	691
iSCSI Target Ports object model	692
IBMTSDS_iSCSIConfigurationService. CreateIPProtocolEndpoint	693
IBMTSDS_iSCSIConfigurationService. ModifyIPProtocolEndpoint	693
IBMTSDS_iSCSIConfigurationService. DeleteIPProtocolEndpoint	694
Creating an IP endpoint	694
Modifying an IP endpoint	694
Deleting an IP endpoint	694
Masking and Mapping profile	695
Masking and Mapping object model	696
IBMTSDS_StorageHardwareIDManagement. CreateStorageHardwareID	697
IBMTSDS_StorageHardwareIDManagement. DeleteStorageHardwareID	697
IBMTSDS_StorageHardwareIDManagement. CreateHardwareIDCollection	698
IBMTSDS_StorageHardwareIDManagement. AddHardwareIDsToCollection	698
IBMTSDS_ControllerConfigurationService. ExposePaths	698
IBMTSDS_ControllerConfigurationService. HidePaths	698
IBMTSDS_ControllerConfigurationService. DeleteProtocolController	699
Sample recipe	699
Indication profile	699
Replication Services profile	700
Replication Services object model	701
Group management	703
IBMTSDS_ReplicationService. CreateGroup	703
IBMTSDS_ReplicationService. DeleteGroup	704
IBMTSDS_ReplicationService. AddMembers	704
IBMTSDS_ReplicationService. RemoveMembers	704
Replication management	704
IBMTSDS_ReplicationService. CreateElementReplica	705
IBMTSDS_ReplicationService. CreateGroupReplica	706
IBMTSDS_ReplicationService. ModifyReplicaSynchronization	706
IBMTSDS_ReplicationService. ModifyListSynchronization	706
IBMTSDS_ReplicationService. GetAvailableTargetElements	706
IBMTSDS_ReplicationService. GetReplicationRelationships	707
IBMTSDS_ReplicationService. GetReplicationRelationshipInstances	707
IBMTSDS_ReplicationService. GetPeerSystems	707
IBMTSDS_ReplicationService. CreateGroupReplicaFromElements	707
Job Control profile	708
Job Control object model	708
Thin provisioning profile	709
Thin Provisioning object model	709
IBMTSDS_PrimordialStoragePoolCapabilities. CreateSetting	710
IBMTSDS_VirtualPoolSetting. modify	710
IBMTSDS_VirtualPoolSetting. delete	710
IBMTSDS_StorageConfigurationService. CreateOrModifyStoragePool	710
Creating a thin provisioning pool	710
Thin provisioning indications	711
Conformance tests	711
Publications	712
IBM XIV Storage System version 11.6.x publications	712
IBM XIV Storage System version 11.5.x publications	713
Warranty, notices, and licensing publications	713
Related publications and websites	714
Glossary	715
IBM XIV 10.2.x publications	719
Release notes publications	720

IBM XIV Storage System documentation

IBM XIV Storage System is a high-end, open, disk-storage system. The XIV system has an innovative grid architecture designed to deliver the highest levels of reliability, performance, scalability and functionality at low overall cost, while eliminating complexity and providing unprecedented ease of management.

Find documentation

[IBM XIV Gen3 11.x](#)

[IBM XIV 10.2.4.x](#)

[Platform and application integration](#)

Learn more

[IBM XIV Storage System website](#)

[Whitepapers](#)

[Redbooks](#)

[Podcasts](#)

Get support

[IBM Support website](#)

[Fixes and downloads](#)

[Service requests and PMRs](#)

Connect with experts

[IBM System Storage education](#)

[IBM Systems Storage forum](#)

[DeveloperWorks technical library](#)

Stay current

[Twitter](#)

.

(C) Copyright IBM Corporation 2008, 2020.

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)

Last updated: 23 November 2018

Note: The listed functional features are supported by using:

- IBM Hyper-Scale Manager management GUI version 5.5.0 or later
- IBM Extended Command-Line Interface (XCLI) utility version 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 version 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](#).

Last updated: 20 July 2017

From Primary/To Secondary	To IBM XIV Gen2	To IBM XIV Gen3	To IBM Spectrum Accelerate	To IBM FlashSystem A9000 version 12.1.0 or later	To IBM FlashSystem A9000R version 12.1.0 or later
From IBM XIV Gen2					
From IBM XIV Gen3					
From IBM Spectrum Accelerate					
From IBM FlashSystem A9000 version 12.1.0 or later				✓	✓
From IBM FlashSystem A9000R version 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](#).

Last updated: 20 July 2017

From Primary/To Secondary	To IBM XIV Gen2	To IBM XIV Gen3	To IBM Spectrum Accelerate	To IBM FlashSystem A9000 version 12.0.1 or later	To IBM FlashSystem A9000R version 12.0.1 or later
From IBM XIV Gen2	✓	✓			
From IBM XIV Gen3	✓	✓	✓		
From IBM Spectrum Accelerate		✓	✓		
From IBM FlashSystem A9000 version 12.0.1 or later				✓	✓
From IBM FlashSystem A9000R version 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](#).

Last updated: 24 November 2017

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

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

Multi-site HA/DR (High Availability & Disaster Recovery) for system software version 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](#).

Last updated: 7 November 2018

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

Multi-site mirroring ("3-way") for system software version 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](#).

Last updated: 1 November 2018

From Primary/To Secondary	To IBM XIV Gen3	To IBM Spectrum Accelerate
From IBM XIV Gen3 version 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](#).

Last updated: 20 March 2018

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

Last updated: 1 November 2018

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 version 12.0.0 or later	To IBM FlashSystem A9000R version 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 version 12.0.0 or later	✓	✓	✓	✓	✓
From IBM FlashSystem A9000R version 12.0.0 or later	✓	✓	✓	✓	✓
From SCSI Storage Array	✓	✓	✓	✓	✓

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](#).

IBM XIV Storage System Gen3 (11.x) documentation

Last updated: 2024-05-22

Welcome to the IBM XIV Storage System Gen3 documentation, where you can find information about the IBM XIV Storage System Gen3 features and operation.

Getting started

- [What's new](#)
- [Release notes](#)
- [Product overview](#)
- [Specifications](#)
- [Planning](#)
- [Overview of XCLI utility](#)
- [PDF publications](#)

Common tasks

- [Administering users](#)
- [Managing via GUI and XCLI commands](#)
- [Configuring systems](#)
- [Compressing data](#)
- [Migrating data](#)
- [Preparing for microcode upgrade](#)

Troubleshooting and support

- [Known issues and workarounds](#)
- [XIV Remote Support Proxy](#)
- [XIV Storage System on IBM Fix Central](#)

Best Practices

- [Redbooks](#)
- [Agreements and license information](#)
These license agreements and program license information documents apply to IBM XIV Storage System.
- [Notices](#)
This information was developed for products and services offered in the U.S.A.
- [Warranty information](#)
The statement of limited warranty specifies the warranty period, type of warranty service, and service level.

Agreements and license information

These license agreements and program license information documents apply to IBM® XIV® Storage System.

These documents are for your reference and viewing convenience. They are normally provided with IBM software in hardcopy form, on a CD, or are embedded in the software in softcopy form for acceptance prior to installation of the software.

Table 1. Product agreements and license information

Title	Order number
<i>International Agreement for Acquisition of Software Maintenance</i>	Z125-6011-04
<i>License Agreement for Machine Code</i>	SC28-6872-02
<i>International Program License Agreement</i>	Z125-3301-14

See the *Agreements and License Information* CD that was shipped with the XIV system for the following documents:

- IBM XIV Storage System License Information
- IBM XIV Storage System License Notices and Information
- IBM XIV Storage System License Supplemental Notices and Information

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

Notices

This information was developed for products and services offered in the U.S.A.

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
Armonk, NY 10504-1785
U.S.A.*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan*

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: 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.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Corporation
Almaden Research
650 Harry Road
Bldg 80, D3-304, Department 277
San Jose, CA 95120-6099
U.S.A.*

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

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

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

- [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.

- [Safety and environmental notices](#)

Review the safety notices, environmental notices, and electronic emission notices for IBM XIV® Storage System before you install and use the product.

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

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Linux® and the Linux logo is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

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

Other product and service names might be trademarks of IBM or other companies.

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 IBM® XIV® Storage System before you install and use the product.

- [Safety notices and labels](#)
Review the safety notices and safety information labels before using this product.
- [Special caution and safety notices](#)
This information describes special safety notices that apply to the IBM XIV Storage System. 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. 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.

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](#).

Special caution and safety notices

This information describes special safety notices that apply to the IBM® XIV® Storage System. 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 IBM approved UL power cables.
- [Leakage current](#)
XIV systems incorporate 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 XIV 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 Performance Standards, 21 CFR Subchapter J, and the international laser safety standard IEC 825-2.

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 radiation performance standards, 21 CFR Subchapter J.

Internationally, use only SFP or GBIC optical transceivers that comply with IEC standard 825-1. 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.

Power cables

Use only IBM® approved UL power cables.

For your safety, IBM provides a power cable with a grounded attachment plug to use with this IBM product. To avoid electrical shock, always use the power cable and plug with a correctly grounded outlet. IBM power cables used in the United States and Canada are listed by Underwriters Laboratories (UL) and certified by the Canadian Standards Association (CSA). For units intended to be operated at 115 volts: Use a UL-listed and CSA-certified cable set consisting of a minimum 18 AWG, Type SVT or SJT, three-conductor cable, a maximum of 15 feet in length and a parallel blade, grounding-type attachment plug rated 15 amperes, 125 volts. For units intended to be operated at 230 volts (U.S. use), use a UL-listed and CSA-certified cable set consisting of a minimum 18 AWG, Type SVT or SJT, three-conductor cable, a maximum of 15 feet in length and a tandem blade, grounding-type attachment plug rated 15 amperes, 250 volts. For units intended to be operated at 230 volts (outside the U.S.), use a cable set with a grounding type attachment plug. The cable set must have the appropriate safety approvals for the country in which the equipment is to be installed. IBM power cables for a specific country or region are usually available only in that country or region.

Connect all power cables to a correctly wired and grounded electrical outlet. Ensure that the outlets supplies correct voltage and phase rotation according to the system rating plate. Ensure that all customer facility outlets are protected with circuit breakers rated at maximum for 30 Amps. The power cable plugs operate as the system main-disconnection method.

Leakage current

XIV® systems incorporate 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 XIV systems. The XIV 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 XIV® system. The customer is responsible for using professional movers or riggers in the case of equipment relocation or disposal.

Attention: Do not tilt the XIV system rack more than 12.5 degrees.

If more clearance is needed for models 281x-114 or 281x-214 racks, the weight-and-height-reduced shipping (feature code 0200) must be ordered to reduce the height of the rack. When this option is ordered, data modules 12 - 15 are shipped separately, and the rack top cover can be removed by the IBM service representative before the XIV system is moved to its final location by professional movers. For more information, see [Weight-and-height reduced shipping feature](#).

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 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-Mitgliedsstaaten und hält die Grenzwerte der EN 55032 Klasse A ein.

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

EN 55032 Klasse A Geräte müssen mit folgendem Warnhinweis versehen werden:

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

Deutschland: Einhaltung des Gesetzes über die elektromagnetische Verträglichkeit von Geräten

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

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

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

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

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

Der verantwortliche Ansprechpartner des Herstellers in der EU ist:

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

Generelle Informationen:

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

Japan Electronics and Information Technology Industries Association (JEITA) Notice

(一社) 電子情報技術産業協会 高調波電流抑制対策実施
要領に基づく定格入力電力値：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情報技術装置です。この装置を家庭環境で使用すると電磁妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

VCCI-A

Korea Notice

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

People's Republic of China Notice

警告

此为A级产品，在生活环境中，该产品可能会造成无线电干扰。在这种情况下，可能需要用户对其干扰采取切实可行的措施。

Russia Notice

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

02/2016

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

12c00790

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

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](#).

What's new

The following table lists the dates and nature of updates to the published information of IBM® XIV® Storage System Gen3 in IBM Documentation since January 2015.

Important: As of June 2022, the IBM XIV Management Tools for XIV Storage System is no longer being maintained. This includes the following : XIV Management Tools 4.8.0.6 and lower, XIV XCLI 4.8.0.6 and lower, and IBM Hyper-Scale Manager v1.9.6 and lower. If these utilities are currently being used, it is important to migrate to the latest IBM Hyper-Scale Manager and XCLI Utility for the most up-to-date security patches. These can both be downloaded from [Fix Central](#). For more information about the XCLI Utility, see [Release notes for IBM XCLI Utility](#) and [Command-line reference](#). For more information about Hyper-Scale Manager, see [Hyper-Scale Manager documentation](#).

Note: The term *documentation* includes the complete set of relevant information deliverables in both PDF and HTML format.

Date	Nature of updates to the published information
18 January 2023	- Second edition release notes for IBM Remote Support Proxy utility version 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.
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.
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 more information, see What's new in 5.5.4. - New edition of the CLI user guide. This edition updates installation information for the XCLI Utility. For more information, see Installing and starting the XCLI. - 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 version 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.
22 July 2021	- Seventh edition release notes for IBM Remote Support Proxy utility version 1.3.2.1. This edition adds PXY-235410 as a new known issue. - For the PDF version of the release notes, see the release notes PDF publications page.
27 December 2020	- Sixth edition release notes for IBM Remote Support Proxy utility version 1.3.2.1. The new edition includes updates in the System requirements section. - For the PDF version of the release notes, see the release notes PDF publications page.
27 July 2020	- 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.
17 March 2020	New edition of the IBM Remote Support Proxy utility version 1.3.2.1 User Guide. The new edition includes updates to the Configuring the Remote Support Proxy service section. For the PDF version of this guide, see the PDF publications page.
29 August 2019	- First edition release notes for IBM XIV Storage System Gen3 version 11.6.2.d. - New edition of the CLI Reference Guide, including new security configuration commands. - For information about what's new in this version, see What's new in version 11.6.2.d. - For the PDF version of the release notes, see the release notes PDF publications page. - For the PDF version of the CLI Reference Guide, see the PDF publications page.
31 July 2019	- Fifth edition release notes for IBM Remote Support Proxy utility version 1.3.2.1. The new edition includes updates in the supported storage systems. - New edition of the IBM Remote Support Proxy utility version 1.3.2.1 User Guide. The new 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.
1 August 2018	First edition release notes for IBM XIV Storage System Gen3 version 11.6.2.c. For the PDF version of the release notes, see the release notes PDF publications page.
1 May 2018	Release of the IBM XIV Storage System Gen3 user documentation set for version 11.6.2.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.
20 March 2018	New edition of the Planning Guide. For the PDF version of the Planning Guide, see the PDF publications page.
21 February 2018	- Fourth edition release notes for IBM Remote Support Proxy utility version 1.3.2.1. The new edition includes updates in the system requirements and supported storage systems. - Third edition release notes for IBM XIV Storage System Gen3 version 11.6.2.a. The following issue has been added to the 11.6.2.a change log: SYS-287670. - For the PDF version of the release notes, see the release notes PDF publications page.
30 January 2018	Third edition release notes for IBM XCLI Utility version 5.2. The new edition includes updates in the "Compatibility" section. For the PDF version of the release notes, see the release notes PDF publications page.
1 November 2017	Second edition release notes for IBM XIV Storage System Gen3 version 11.6.2.a. The following issue has been added to the 11.6.2.a known issues: - SYS-297348 - For the PDF version of the release notes, see the release notes PDF publications page.
15 October 2017	New edition of the CLI Reference Guide, including an expanded description of the <code>COMPRESSED_VOLUME_MARGINAL_COMPRESSION</code> event in the <code>vol_transform_create</code> command. For the PDF version of the CLI Reference Guide, see the PDF publications page.
11 August 2017	New edition of the Open API Guide, including a new note added regarding the use of special characters in user names and passwords. - For the HTML version of the updated Open API Guide topic, see Configuring the CIM agent. - For the PDF version of the Open API Guide, see the PDF publications page.

Date	Nature of updates to the published information
2 August 2017	Third edition release notes for IBM Remote Support Proxy utility version 1.3.2.1. The new edition includes updates in the system requirements and supported storage systems. For the PDF version of the release notes, see the release notes PDF publications page.
1 August 2017	Second edition release notes for IBM XIV Management Tools version 4.8.0.6. The What's new in version 4.8.0.6 section has been updated.
17 July 2017	Release of the IBM XIV Storage System Gen3 user documentation set for version 11.6.2.a, including a new edition of the IBM Hyper-Scale Manager release notes. - 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
16 January 2017	- Second edition release notes for IBM XIV Storage System Gen3 version 11.6.2. The following issues have been added to the 11.6.2 change log: - SYS-281678 (fixed issue) - SYS-281984 (fixed issue) - SYS-276398 (fixed issue) - SYS-283346 (enhancement) - Planning information update - Added IBM Hyper-Scale Manager 5.1.1 documentation in HTML web-page format, in addition to PDF.
12 December 2016	- IBM XIV Storage System Gen3 11.6.2 documentation - IBM XIV Management Tools 4.8.0.4 release notes
8 November 2016	- IBM Hyper-Scale Manager 5.1 documentation (PDF only) - IBM Remote Support Proxy 1.3.2.1 release notes
23 September 2016	Third edition release notes for IBM XIV Storage System Gen3 version 11.6.1.a.
26 July 2016	Planning information updates
14 July 2016	IBM XIV Management Tools 4.8.0.3 release notes
28 April 2016	IBM XIV Management Tools 4.8.0.2 release notes
30 March 2016	- IBM XIV Storage System Gen3 11.6.1.a release notes - IBM XIV Management Tools 4.8.0.1 release notes
22 December 2015	- IBM XIV Management Tools 4.8 documentation - IBM Hyper-Scale Manager 1.9 documentation
4 December 2015	- IBM XIV Storage System Gen3 11.6.1 documentation - Planning information for the new 281x-314 XIV model
27 September 2015	- IBM XIV Management Tools 4.7 documentation - IBM Hyper-Scale Manager 1.8 documentation
25 September 2015	IBM XIV Storage System Gen3 11.4.2.e release notes
27 August 2015	IBM XIV Storage System Gen3 11.5.2 release notes
25 June 2015	- IBM XIV Management Tools 4.6.0.1 release notes - IBM XIV Management Tools 4.5.0.2 release notes
4 June 2015	- New organization of the XIV Gen3 main topics on the table of contents - IBM XIV Storage System Gen3 11.6 documentation - IBM XIV Management Tools 4.6 documentation - IBM Hyper-Scale Manager 1.7 documentation
8 May 2015	IBM XIV Management Tools 4.5.0.1 release notes
30 April 2015	IBM XIV Storage System Gen3 11.4.2.d release notes
2 April 2015	IBM XIV Storage System Gen3 11.5.1.b release notes
13 March 2015	- IBM XIV Management Tools 4.5 documentation - IBM Hyper-Scale Manager 1.6 documentation
18 February 2015	- IBM XIV Storage System Gen3 release notes for the following versions: 11.5.1.a 11.4.2.c - IBM XIV Management Tools 4.4.1.1 release notes - IBM Hyper-Scale Manager 1.5.1.33 release notes

Date	Nature of updates to the published information
12 February 2015	IBM XIV Storage System Gen3 11.5.1 documentation

Release notes

Refer to the following sections for release information for the latest releases of IBM® XIV® Storage System Gen3, IBM Hyper-Scale Manager, IBM XIV Management Tools, IBM Extended Command-Line Interface (XCLI) Utility, and IBM Remote Support Proxy, as well as access to PDF versions of current and previous releases.

- [Release notes for IBM XIV Storage System](#)
The following release information is available for version 11.6.2.d of IBM XIV Storage System.
- [Release notes for IBM Hyper-Scale Manager](#)
For release information about IBM Hyper-Scale Manager, see the 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.
- [Release notes for IBM XIV Management Tools](#)
The following release information is available for version 4.8.0.6 of IBM XIV Management Tools.
- [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 notes PDF publications](#)
Use this page to access PDF versions of release notes for IBM XIV Storage System Gen3, IBM XIV Management Tools, IBM Extended Command-Line Interface (XCLI) Utility, and IBM Remote Support Proxy.

Release notes for IBM XIV Storage System

The following release information is available for version 11.6.2.d of IBM® XIV® Storage System.

- [What's new in version 11.6.2.d](#)
Software version 11.6.2.d enhances the operation of existing functional features and resolves miscellaneous issues.
- [Hot upgrade from earlier versions](#)
Version 11.6.2.d supports direct hot upgrade from the following XIV Gen3 software versions.
- [Management and solutions](#)
IBM XIV Gen3 version 11.6.2.d can be managed and monitored with the following software management tools and cloud storage solutions.
- [Change log](#)
This change log summarizes the enhancements and fixes in the different 11.6.x code level versions of IBM XIV Storage System Gen3.
- [Fixes and features of earlier versions](#)
IBM XIV Storage System Gen3 version 11.6.2.d includes all the features and fixes that were included in previous 11.x.x versions.
- [Limitations](#)
The following functional limitations apply to IBM XIV Storage System Gen3 and may remain in effect in future XIV software versions.
- [Known issues](#)
This section details the known issues in IBM XIV Storage System Gen3, along with possible solutions or workarounds (if available).

What's new in version 11.6.2.d

Software version 11.6.2.d enhances the operation of existing functional features and resolves miscellaneous issues.

General availability date: 29 August 2019

Added security configuration CLI commands

Added **protocol_config_set** XCLI command to configure a higher than default minimum of the SSL/TLS protocol (SSLv3). **protocol_config_list** XCLI command was added in order to list the current configuration parameters.

Resolved issues

For information about the resolved issue in version 11.6.2.d, see [Change log](#).

Important: Unless using IBM Hyper-Scale Manager version 5.5.2 or later to monitor and manage the storage system, XIV Gen3 version 11.6.2.d requires mandatory upgrade of the XIV Management Tools to version 4.8.0.6 or later. For more information, see [Management and solutions](#).

Hot upgrade from earlier versions

Version 11.6.2.d supports direct hot upgrade from the following XIV® Gen3 software versions.

- 11.6.1, 11.6.1.a, 11.6.2, 11.6.2.a, 11.6.2.b, 11.6.2.c

Other 11.x.x versions can be upgraded to 11.6.2.d by applying consecutive hot upgrades, depending on the version in use.

Note: These hot upgrade paths are valid for the time of the general availability (GA) date of version 11.6.2.d. However, hot upgrade from additional earlier versions could be supported later on. For up-to-date information, contact IBM® Support.

The upgrade must be performed only by an authorized IBM service technician.

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.

Before initiating the upgrade, some preparation steps are required. For more information, see 'Non-disruptive code load' in the IBM XIV Gen3 Product Overview publication.

Management and solutions

IBM® XIV® Gen3 version 11.6.2.d can be managed and monitored with the following software management tools and cloud storage solutions.

- IBM Hyper-Scale Manager version 5.5.2 or later – provides an advanced and brand-new web-based graphical user interface (GUI) from which one or more XIV systems (and any other IBM Spectrum Accelerate Family system) can be managed and monitored in real time from a web browser.
- IBM XIV Management Tools version 4.8.0.6 or later
- IBM XIV Mobile Dashboard (for monitoring only) version 1.4.0.1 or later

Important: Although earlier management tool versions may be able to detect XIV systems of version 11.6.2.d, the latest XIV functional features and/or Management Tools security updates are available only in the versions specified above.

IBM XIV Gen3 can also be managed from its command-line interface (CLI) or through its RESTful application programming interfaces (APIs).

XIV customers are entitled to free use of these management tools, which can be downloaded at any time from the [IBM Fix Central website](#).

The Mobile Dashboard application can be downloaded and installed at any time and from [Apple Store](#) (iOS version) or [Google Play](#) (Android version).

IBM Spectrum Control Base Edition

IBM XIV Gen3 version 11.6.2.d is supported by IBM Spectrum Control Base Edition version 3.0.3 or later, which provides the following VMware platform integration tools.

- IBM Storage Provider for VMware VASA
- IBM Storage Enhancements for VMware vSphere Web Client
- IBM Storage Plug-in for VMware vCenter Orchestrator
- IBM Storage Management Pack for VMware vCenter Operations Manager

For more information, see [IBM Spectrum Control Base Edition archived publications](#).

XIV customers are entitled to free use of IBM Spectrum Control Base Edition, which can be obtained at any time from either the IBM Hyper-Scale Manager 5.x.x software package or from the [IBM Fix Central website](#).

Available cloud storage solutions

IBM XIV Storage System Gen3 version 11.6.2.d 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 Host Attachment Kit	2.6.0 or later	Available for IBM AIX®, HP Unix, Linux®, Oracle Solaris, and Microsoft Windows Server
IBM XIV Provider for Microsoft Windows Volume Shadow Copy Service (VSS)	2.8.0.1 or later	
IBM XIV Storage Replication Adapter (SRA)	2.3.0 or later	For use with VMware vCenter Site Recovery Manager
IBM Storage Driver for OpenStack	1.8.0 or later	Version 1.8.0 supports only the OpenStack Newton release
IBM Storage Enabler for Windows Failover Clustering	1.2.0 or later	Version 1.2.0 does not support compressed volumes (through the IBM Real-time Compression feature)
IBM Storage Management Pack for Microsoft System Center Operations Manager	2.6.0 or later	

XIV customers are entitled to free use of the IBM cloud storage solutions, which can be downloaded at any time from the [IBM Fix Central website](#).

For more information about the available cloud storage solutions for XIV, see [Platform and application integration](#) on IBM Documentation .

Change log

This change log summarizes the enhancements and fixes in the different 11.6.x code level versions of IBM® XIV® Storage System Gen3.

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.

- [11.6.2.d \(August 2019\)](#)
Version 11.6.2.d of IBM XIV Storage System Gen3 resolves the following issues.
- [11.6.2.c \(August 2018\)](#)
Version 11.6.2.c of IBM XIV Storage System Gen3 resolved the following issue.
- [11.6.2.b \(May 2018\)](#)
IBM XIV Storage System Gen3 11.6.2.b enhanced the operation of the following functional features and resolved the following issues.
- [11.6.2.a \(July 2017\)](#)
Version 11.6.2.a of IBM XIV Storage System Gen3 resolved the following issues.
- [11.6.2 \(December 2016\)](#)
Version 11.6.2 of IBM XIV Storage System Gen3 resolved the following issues.
- [11.6.1.a \(March 2016\)](#)
Version 11.6.1.a of IBM XIV Storage System Gen3 resolved the following issues.
- [11.6.1 \(December 2015\)](#)
Version 11.6.1 of IBM XIV Storage System Gen3 introduced Turbo Compression in the new 281x-314 XIV hardware model, as well as compression enhancements for all XIV Gen3 models.
- [11.6.0.b \(June 2015\)](#)
Version 11.6.0.b of IBM XIV Storage System Gen3 was the first major release of the 11.6 code level that introduced the following primary features.

11.6.2.d (August 2019)

Version 11.6.2.d of IBM® XIV® Storage System Gen3 resolves the following issues.

Ticket ID	Description
SYS-313334	Fixed: In rare conditions a bug in the microcode might lead to undetected corruption (a range of 1MB returning old data instead of zeros), on the Secondary volume of any type of mirroring relationship: synchronous mirroring, asynchronous mirroring or 3 way mirroring. Severity: HIPER Affected versions: 11.2.0 or later
SYS-276739	Fixed: In rare cases following specific software failures on the secondary system, when using synchronous mirroring, the host connected to the primary system may encounter loss of access to that volume. Severity: High Impact Affected versions: 11.0.0 or later
SYS-287046	Fixed: In rare cases, when 512 LUNs are mapped to a single FC port and then are being rescanned by a host, the rescan may fail, leading to loss of access of the relevant host. Severity: High Impact Affected versions: 11.0.0 or later
SYS-310712	Fixed: In rare cases, when a volume size is not a full multiple of a partition size and the last partition is allocated on the secondary volume and is comprised entirely of 0s, a cache node failure may occur. In extremely rare cases, when two such failures occur in a short period of time, before rebuild is complete, this could result in data loss, due to the loss of both data copies. Severity: High Impact Affected versions: 11.6.2.a or later
SYS-314905	Fixed: Security scans report security vulnerability issues of management IP addresses. 3DES ciphers for XCLI connections (TCP port 7778), are not disabled by default. Severity: High Impact Affected versions: 11.6.2.a 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: 11.6.2.d or later

System and command-line interface (CLI) enhancement that was been added in version 11.6.2.d:

Ticket ID	Description
SYS-314695	Enhancement: To allow the configuration of a higher than default minimum of the SSL/TLS protocol (SSLv3), the protocol_config_set XCLI command was added. For more information, see the description of the protocol_config_set XCLI command in the IBM XIV Storage System Command-Line Interface (CLI) Reference Guide.

11.6.2.c (August 2018)

Version 11.6.2.c of IBM® XIV® Storage System Gen3 resolved the following issue.

Ticket ID	Description

Ticket ID	Description
SYS-306775	Fixed: When an iSCSI command with an invalid data length is sent from a host to a target port, this host, as well as other connected hosts, might lose access to the target port. If the host continues to send the same command, additional target ports may fail, which could ultimately lead to a loss of access. Severity: High Impact Affected versions: 11.0.0 or later

11.6.2.b (May 2018)

IBM® XIV® Storage System Gen3 11.6.2.b enhanced the operation of the following functional features and resolved the following issues.

Ticket ID	Description
SYS-297348	Fixed: During a rebuild process, when medium errors are discovered, the system automatically fixes them. If all the following conditions are met in the course of this operation, an undetected data corruption will occur: • The LBA subject to medium error belongs to a non-compressed volume • The volume is larger than 2TB • The LBA is located beyond 2TB If the related volume is compressed, a medium error will be returned for the affected location, or the volume may go offline. In either case, volumes whose size does not exceed 2TB are not affected. Severity: HIPER Affected versions: 11.2.0 or later
SYS-305463	Fixed: The primary volume in synchronous or asynchronous mirroring, that was restored from a snapshot or was involved in a vol_copy operation at any time, and later also went through offline mirror initialization, role change, or re-synchronization while the system was redistributing data, could suffer from undetected data corruption that is manifested by reading zeros instead of the correct data. In addition, deletion of such volume may lead to loss of access to all the hosts. If such data corruption occurs - regardless of whether the corrupted data was actually read - the issue will be eventually automatically detected by an internal periodical scrub process. Then, if the system is connected to Proactive Support (IBM Call Home), IBM will be automatically notified, and an IBM support person will be engaged to address the issue. Severity: HIPER Affected versions: 11.0.0 or later
SYS-305488	Fixed: The Estimated Sync Time field does not appear in the mirror_list CLI command output even when using the -f option. Severity: Service Affected versions: 11.0.0 or later

11.6.2.a (July 2017)

Version 11.6.2.a of IBM® XIV® Storage System Gen3 resolved the following issues.

Ticket ID	Description
SYS-283435	Fixed: In extremely rare cases, a read operation might be incomplete, exposing the storage system to undetected data corruption. To address this extremely rare situation, an advanced mechanism was introduced, to detect and to protect from such issues. Severity: HIPER Affected versions: 11.0.0 or later
SYS-285605	Fixed: In extremely rare cases, a snapshot of the primary volume used in the asynchronous mirroring might be incomplete. As a result, the data in the secondary volume might be inconsistent with the data on the primary volume. Severity: HIPER Affected versions: 11.0.0 or later
SYS-275412	Fixed: In rare cases, an SSD failure could result in the entire module failure. Severity: High Impact Affected versions: 11.6.2 or later
SYS-282590	Fixed: Performance might deteriorate, when a remote target system in asynchronous mirroring has 6TB disk drives. IBM Technote S1010409, Impact of large sector disk drive on IBM XIV Gen3 mirroring, provides important guidance for using IBM XIV Gen3 storage systems with 6TB disk drives in asynchronous mirroring. Severity: High Impact Affected versions: 11.5.1 or later

Ticket ID	Description
SYS-283154	Fixed: Changing role for a subordinate volume in asynchronous mirroring, while I/Os are being sent to primary volume, might cause a cache node on the original master to stop. This invokes a rebuild operation for all disks in the relevant module. Severity: High Impact Affected versions: 11.5.0 or later
SYS-283199	Fixed: Common OpenSSL security issues. For more information about these security issues, see CVE-2017-2177 , CVE-2017-2178 , CVE-2017-2179 , CVE-2017-2180 , CVE-2017-2181 , CVE-2017-2182 , CVE-2017-2183 , CVE-2017-6302 , CVE-2017-6304 , CVE-2017-6306 on the Common Vulnerabilities and Exposures (CVE) information website Severity: High Impact Affected versions: 11.6.0 or later
SYS-284333	Fixed: In very rare cases, a software deadlock condition may result in a loss of access. Severity: High Impact Affected versions: 11.2.0 or later
SYS-284393	Fixed: A storage system fails to connect to a key server after the system power-up, if the server uses a malformed certificate. Severity: High Impact Affected versions: 11.4.0 or later
SYS-286311	Fixed: Due to an error in the SCSI reservation mechanism, deletion of a host that had a previous reservation could cause management operations to be blocked. This results in I/Os that involve SCSI reservation actions to fail. Severity: High Impact Affected versions: 11.0.0 or later
SYS-286446	Fixed: Common security issues. For more information about these security issues, see CVE-2017-3731 , CVE-2016-8610 , CVE-2016-9131 , CVE-2016-9147 , CVE-2016-9444 , CVE-2016-2776 on the Common Vulnerabilities and Exposures (CVE) information website Severity: High Impact Affected versions: • 11.5.0 or later (CVE-2016-9131, CVE-2016-9147, CVE-2016-9444, CVE-2017-3731, CVE-2016-8610) • 11.6.2 or later (CVE-2016-2776)
SYS-287670	Fixed: Common security issues. For more information about these security issues, see CVE-2016-2775 , CVE-2017-3136 , CVE-2017-3137 , and CVE-2017-3138 on the Common Vulnerabilities and Exposures (CVE) information website . Severity: High Impact Affected versions: 11.6.2 or later
SYS-284698	Fixed: In very rare cases, internal communication malfunction might prevent data rebuild from completion. Severity: Moderate Affected versions: 11.2.0 or later
SYS-285457	Fixed: Mismatch between the reported and actual volume suitability for compression, as notified by the COMPRESSED_VOLUME_MARGINAL_COMPRESSION event. Severity: Service Affected versions: 11.6.0 or later

System and command-line interface (CLI) enhancements that have been added in version 11.6.2.a:

Ticket ID	Description
SYS-269163	Enhancement: To improve certificate management by key server gateways, the following certificate-related values have been changed: • Increase total PEM size of certificate chains up to 15360 characters. • Allow unlimited number of certificates per chain.
SYS-283768	Enhancement: Improved functionality of the Call Home service by: • Allowing storage admins to view notification rule, destination and SMTP gateway lists. • Allowing storage admins to update configuration parameters of SMTP gateways. • Making the SMTP_GATEWAY_FAILED and SMTP_GATEWAY_VIA_NODE_FAILED events visible to administrators.
SYS-285999	Enhancement: Improved volume compression process by introducing the force_on_marginal_ratio parameter for the vol_transform_create command. This parameter prevents compression on volumes, if expected compression ratio is sub-optimal. For more information, see the description of the vol_transform_create command in the IBM XIV Storage System Command-Line Interface (CLI) Reference Guide.
SYS-266624	Enhancement: Improved XIV Support access by enabling multiple source IP addresses, using the xiv_support_enable command. For more information, see the description of the xiv_support_enable command in the IBM XIV Storage System Command-Line Interface (CLI) Reference Guide.

11.6.2 (December 2016)

Ticket ID	Description
SYS-276312	Fixed: In rare cases when using synchronous or asynchronous mirroring, an I/O operation might not be recorded correctly on the slave system when a very small amount of data needs to be synchronized. This will be detected and fixed automatically by the periodic scrub process. Severity: HIPER Affected versions: 11.0.0 or later
SYS-275731	Fixed: Common OpenSSL security issues. For more information about these security issues, see CVE-2015-3194, CVE-2015-3195, CVE-2015-3196, CVE-2016-2107, CVE-2016-2108, CVE-2016-2109, and CVE-2016-2176 on the Common Vulnerabilities and Exposures (CVE) information website . Severity: High Impact Affected versions: 11.0.0 or later
SYS-275992	Fixed: Cache nodes of storage systems with compressed volumes might fail during an upgrade operation. Note: This does not occur when upgrading to software version 11.6.2. Severity: High Impact Affected versions: 11.6.0 or later
SYS-276233	Fixed: A VVol-based virtual machine, which includes a volume with snapshots, that is migrated within the same IBM XIV storage system using VMware Storage vMotion, is exposed to a potential undetected data corruption issue. A read from a volume or snapshot of a migrated VVol-based virtual machine may return incorrect data. Severity: High Impact Affected versions: 11.5.0 or later
SYS-276688	Fixed: Common OpenSSH security issues. For more information about these security issues, see CVE-2016-0777 and CVE-2016-0778 on the Common Vulnerabilities and Exposures (CVE) information website . Severity: High Impact Affected versions: 11.0.0 or later
SYS-277077	Fixed: The ‘switch role’ mirror operation for a consistency group is not blocked if any of the volumes on the previous master remained mapped. This would cause a volume that is now a slave to remain mapped (as opposed to the requirement for unmapping the slave volume) and the data on this volume might become corrupted. Note: To prevent this issue from occurring, this ‘switch role’ operation is blocked starting from version 11.6.2. For more information, see Limitations. Severity: High Impact Affected versions: 11.6.0 or later
SYS-278540	Fixed: In very rare cases that involve timing of user read/write combined with internal system procedure, a user read I/O operation might not be acknowledged and the compression node that is handling the I/O operation might fail. The other compression nodes continue to work as usual and their I/O operations continue to be served. The failed compression node revival must be performed by IBM Support. Severity: High Impact Affected versions: 11.6.0 or later
SYS-281678	Fixed: In rare and specific cases, the asynchronous remote mirroring sync job might get stuck with a single partition left without synchronization. Severity: High Impact Affected versions: 11.4.0 or later
SYS-281984	Fixed: In rare cases, a particular software error might lead to InfiniBand interconnect communication error between one of the interface nodes and one of the cache nodes, leading to host disconnection. Severity: High Impact Affected versions: 11.5.1 or later
SYS-266936	Fixed: When starting and completing a transformation from compressed state to uncompressed state or vice versa, Unit Attention is not raised for all hosts that are mapped to the transformed volume. Severity: Moderate Affected versions: 11.6.0 or later
SYS-276859	Fixed: The XCOPY (Extended Copy) SCSI command, which is part of the VMware vSphere API for Array Integration (VAAI) specification, is not supported with compressed volumes. This impacts the speed of VMware environment copy operations by automatically falling back to a host-based copy. The impacted copy operations are virtual machine cloning, virtual machine creation from a template, and vSphere Storage vMotion, within the same storage array. Severity: Moderate Affected versions: 11.6.0 or later
SYS-276398	Fixed: The event saver node that is responsible for sending events might fail when handling a large number of events. This has minimal impact on the system as there is an event saver instance on each of the other XIV modules. Severity: Service Affected versions: 11.6.0.b or later

Ticket ID	Description
SYS-272799	Fixed: Media errors on the disk might cause the scrub process to stop running. Severity: Service Affected versions: 11.5.0.a or later
SYS-276062	Fixed: Soft and hard size values appear incorrectly in the POOL_RESIZE_THIN event. The value of the soft size is displayed for the hard size, and vice versa. Severity: Service Affected versions: 11.6.0.b or later
SYS-280256	Fixed: In some cases when a mirroring relationship is deleted, the management GUI might display a blank page. Severity: Low Affected versions: 11.0.0 or later

Command-line interface (CLI) enhancement that has been added in version 11.6.2:

Ticket ID	Description
SYS-283346	Enhancement: Added volume listing filtering by World Wide Name (WWN). The wwn parameter can be used with the vol_list command in the following format: vol_list wwn=<wwn number> . The WWN number of all volumes can be obtained by issuing the following command: vol_list -f name,wwn

11.6.1.a (March 2016)

Version 11.6.1.a of IBM® XIV® Storage System Gen3 resolved the following issues.

Ticket ID	Description
SYS-277205	Fixed: In rare cases where a module fails and another module boots slowly during a microcode hot upgrade, the cache nodes might fail. Severity: High Impact Affected versions: 11.5.1 or later
SYS-277761	Fixed: GNU C Library vulnerability issue (not specific to XIV systems): GNU C Library (glibc) is vulnerable to a stack-based buffer overflow, caused by improper bounds checking by the NSS DNS backend for the getaddrinfo() function when performing dual A/AAAA DNS queries. By sending a specially crafted DNS response, a remote attacker could overflow a buffer and execute arbitrary code on the system, or cause the application to crash. For more information about this security issue, see CVE-2015-7547 on the Common Vulnerabilities and Exposures (CVE) information website . Severity: High Impact Affected versions: 11.0.0 or later

11.6.1 (December 2015)

Version 11.6.1 of IBM® XIV® Storage System Gen3 introduced Turbo Compression in the new 281x-314 XIV hardware model, as well as compression enhancements for all XIV Gen3 models.

Fixes in version 11.6.1:

Ticket ID	Description
SYS-252858	Fixed: In rare cases, when a data redistribution process runs in parallel with a synchronous or asynchronous mirroring synchronization process, host disconnection might occur. Severity: High Impact Affected versions: 11.0.0 or later
SYS-271569	Fixed: In rare cases, a failure of an SSD cache drive might result in host disconnection. Severity: High Impact Affected versions: 11.5.0 or later
SYS-273223	Fixed: In rare cases, a system with an asynchronous mirror relation might fail if the target connectivity is down and a new asynchronous mirror is created with the Offline init option selected. Severity: High Impact Affected versions: 11.0.0 or later
SYS-275362	Fixed: In extremely rare cases, a race condition might cause data loss in a certain compressed volume. As a result, the volume goes offline. Severity: High Impact Affected versions: 11.6.0 or later
SYS-271892	Fixed: When setting Microsoft Azure Site Recovery protection between two sites using storage area network (SAN) replication, if a non-2810 XIV machine type is selected as the storage provider on either the local site or remote site, protection operations fail. Severity: Moderate Affected versions: 11.6.0 or later

Ticket ID	Description
SYS-273769	Fixed: In rare cases, a non-critical thread might enter an endless loop that affects performance. The impact on performance could range from negligible to significant, depending on the system work load. Severity: Moderate Affected versions: 11.0.0 or later
SYS-274270	Fixed: strongSwan vulnerability issue (not specific to XIV systems): strongSwan could allow a remote authenticated attacker to obtain sensitive information, caused by an error in IKEv2 connections related to server authentication with a certificate and EAP or pre-shared keys. An attacker could exploit this vulnerability to obtain user credentials and other sensitive information. For more information about this security issue, see CVE-2015-4171 on the Common Vulnerabilities and Exposures (CVE) information website . Severity: Moderate Affected versions: 11.0.0 or later
SYS-274413	Fixed: Net-SNMP vulnerability issue (not specific to XIV systems): Net-SNMP is vulnerable to a denial of service, caused by incompletely parsed VarBind variables being left in the list of variables by the <code>snmp_pdu_parse()</code> function. A remote attacker could exploit this vulnerability to cause the application to crash or possibly execute arbitrary code on the system. For more information about this security issue, see CVE-2015-5321 on the Common Vulnerabilities and Exposures (CVE) information website . Severity: Moderate Affected versions: 11.0.0 or later
SYS-273234	Fixed: In a rare cases, a memory corruption might cause a certain software component to fail, resulting in a slight decrease in performance of compressed volumes. Severity: Service Affected versions: 11.6.0 or later
SYS-273704	Fixed: In some cases, management systems that use SMI-S for controlling the XIV system are unable to control it. Severity: Service Affected versions: 11.0.0
SYS-273895	Fixed: When an LDAP directory user tries to accept the license for a new system or after a microcode upgrade that includes a new license, the operation fails with the error: <code>ILLEGAL_NAME</code> Severity: Service Affected versions: 11.0.0
SYS-273902	Fixed: Users with Application Admin role might be unable to create a snapshot of a consistency group. This occurs if the consistency group contains one or more volumes that are not mapped to a host that is associated with a user group that contains these users. Severity: Service Affected versions: 11.1.0 or later
SYS-274349	Fixed: Some systems might stop sending heartbeat info to IBM Support. As a result, IBM Support might not have complete information when handling arising issues. Severity: Service Affected versions: 11.6.0 or later

Security improvement in version 11.6.1:

Ticket ID	Description
SYS-273930	Since version 11.5.0, users with Storage Administrator role were able to create Security Administrator users, while users with Security Administrator role were able to change the password of the default administrator. Starting from version 11.6.1: 1. Users with Security Administrator role cannot change password of the default administrator. 2. Only the default administrator can create, rename, change the password of, or delete Security Administrator accounts.

11.6.0.b (June 2015)

Version 11.6.0.b of IBM® XIV® Storage System Gen3 was the first major release of the 11.6 code level that introduced the following primary features.

- IBM Real-time Compression
- Support for integration with Microsoft Azure Site Recovery
- Automatic shutdown in extreme overheating conditions

Fixes in version 11.6.0.b:

Ticket ID	Description
SYS-269182	Fixed: In extremely rare situations on systems with SSD caching enabled, when two read operations are sent by the host to the same address, the second read operation might return incorrect data from the cache, which could result in undetected data corruption. Subsequent reads to the same address (after the data is no longer in the cache) returns the correct data. Severity: High Impact Affected versions: 11.1.0 or later

Ticket ID	Description
SYS-252743	Fixed: When entering several certificates to the <code>certificate</code> parameter, <code>ipsec_connection_*</code> XCLI commands processes only the first certificate. Severity: Service Affected versions: 11.2.0 or later
SYS-265572	Fixed: If any of the parameters of <code>pool_change_config</code> are invalid, the command might start running and then fail halfway through. Severity: Service Affected versions: 11.5.0 or later
SYS-266212	Fixed: If a solid-state drive fails during the encryption process, the <code>SSD_HAS_FAILED_WHILE_ENABLING_ENCRYPTION</code> event might not be issued. Severity: Service Affected versions: 11.5.0 or later
SYS-266220	Fixed: The <code>format=full</code> parameter on the XCLI <code>help</code> command is not supported. Affected versions: 11.5.0 or later Severity: Service
SYS-266221	Fixed: The <code>ldap_user_list role=storageadmin</code> XCLI command does not list domain administrators. Severity: Service Affected versions: 11.5.0 or later
SYS-266476	Fixed: When using the multi-tenancy feature in LDAP authentication mode, a domain with which users are currently associated cannot be deleted. Severity: Service Affected versions: 11.5.0 or later
SYS-266853	Fixed: In the XCLI help, the descriptions of the <code>domain_manage</code> and <code>domain_move_gp</code> commands are incorrect. Severity: Service Affected versions: 11.5.0 or later
SYS-267911	Fixed: IBM Hyper-Scale Manager supports LDAP UPN mode only when there are no domains configured on the system. If domains are configured on the system, LDAP UPN mode is not supported. Severity: Service Affected versions: 11.5.0 or later

Fixes and features of earlier versions

IBM® XIV® Storage System Gen3 version 11.6.2.d includes all the features and fixes that were included in previous 11.x.x versions.

To obtain information regarding previously introduced XIV Gen3 features or issues that were resolved in previous versions, see [Release notes PDF publications](#)

Limitations

The following functional limitations apply to IBM® XIV® Storage System Gen3 and may remain in effect in future XIV software versions.

Note: For information about specification-related functional boundaries, such as maximum number of volumes and hosts per storage system, refer to the IBM XIV Storage System Gen3 specification sheet document, which can be provided to IBM customers under non-disclosure agreement (NDA).

Limitations with IBM Real-time Compression

The following limitations apply to the IBM Real-time Compression feature when used on XIV Gen3 systems:

- The minimum volume size that can be compressed is 51 Gigabytes (GB). The compression option is blocked for smaller volumes.
- The maximum volume size that can be compressed is 10 Terabytes (TB). The compression option is blocked for larger volumes.
- Up to 1024 volumes can be compressed on a single XIV system.
- Only thin-provisioned volumes can be compressed.
- Shrinking a compressed volume is not supported.
- Snapshots cannot be converted from compressed to uncompressed and vice versa. Snapshots that already existed before a volume was converted from non-compressed to compressed are not converted and are not available with the converted volume.
- When migrating volumes using IBM Hyper-Scale Mobility, compression or decompression operations can only begin after the migration is complete.
- Locked volumes must be unlocked before compression or decompression can begin.
- In a mirroring relation, both the source and target volumes must be compressed, or both must be uncompressed.
- Mapping a host to a compressed volume on a slave node in a mirrored system is not permitted.
- To convert a mirrored volume from uncompressed to compressed (or vice versa), the mirroring relationship for that volume must first be removed and then recreated after the conversion.
- When two-way or three-way mirroring is used, compressed volumes on the secondary or tertiary system cannot be mapped to hosts, except when these volumes become primary volumes after a failover operation. Before any failover operation (including failover testing) compressed primary (master) volumes must first be unmapped from their hosts.
- IBM PowerHA® Enterprise Edition is not supported.

- The latency of synchronously mirrored compressed volumes could be higher than that of synchronously mirrored uncompressed volumes. As a result, if the distance between the production and recovery sites is extremely large, host disconnections might occur. It is therefore recommended to test the latency on a test volume before using synchronous mirroring with compressed volumes in the production environment.
- The system continuously assesses potential saving ratio at a pace of approximately 20 volumes per hour. Until the initial assessment of any non-compressed volume is complete, its potential saving ratio is not displayed.
- The calculated potential saving percentage of a decompressed volume (converted from compressed to non-compressed) could be different than the saving percentage that was originally calculated for the same volume prior to its compression. This occurs if most of the data on the volume comprises zeros, while other data is sparsely distributed.
- VMware vSphere Virtual Volumes (VVOs) are not supported with compressed volumes.
- Data migration CLI commands (**dm_activate**, **dm_deactivate**, **dm_define**, **dm_delete**, **dm_list**, **dm_test**) are not supported with compressed volumes.

Limitations with IBM Hyper-Scale Mobility

The following limitations apply to the IBM Hyper-Scale Mobility (volume migration) feature when used with XIV Gen3 systems:

- Can be used only with systems that are located within the same datacenter.
- Cannot be used on a volume that belongs to a consistency group. The volume must first be removed from the consistency group.
- Does not apply to volume snapshots. When the migration is changed to proxy state, all snapshots of the pertinent migrated volume are deleted automatically.
- Does not migrate volume mirroring relationships. Before a volume can be moved, its mirroring relationship must be deleted.

Known issues

This section details the known issues in IBM® XIV® Storage System Gen3, 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 11.6.2.d or earlier versions. As long as a newer version has not yet been released, a newer release notes edition for version 11.6.2.d 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 11.6.2.d will no longer be updated. Accordingly, check the release notes of the newer version to learn whether any newly discovered issues affect version 11.6.2.d or whether the newer version resolves any of the issues listed below.			
Ticket ID	Severity	Affected versions	Description
SYS-260642	High Impact	11.4.0 or later	The following XCLI commands require a user name as a parameter, but appear to run successfully even if the specified user does not exist: • encrypt_recovery_key_generate • encrypt_keyserver_rekey • encrypt_recovery_key_rekey Workaround: Before running these commands, use the user_list XCLI command to make sure that the user exists.
SYS-267506	High Impact	11.5.1 or later	On systems with 6 TB drives, a single restore of a volume from a snapshot or a consistency group (CG) cannot be handled when the total hard size of the volumes involved (including all their snapshots) exceeds 90% of the system's entire usable capacity. Workaround: Avoid creating a single volume or a single CG with a total hard size that exceeds 50% of the system usable capacity.
SYS-253004	Moderate	11.0.0 or later	Resizing a pool (reducing its volume capacity) does not trigger the STORAGE_POOL_VOLUME_USAGE_INCREASED event when crossing the preset threshold. Workaround: No workaround is currently available. Accordingly, use caution when resizing down a storage pool.
SYS-264917	Moderate	11.5.0 or later	In a multi-site mirroring relation, if the destination remains connected to the source after the source is disconnected from the host and the secondary source, establishing an asynchronous mirror from the secondary source to the destination requires a full resynchronization. Workaround: No workaround is currently available.
SYS-265764	Moderate	11.0.0 or later	When changing a volume to read only, the system does not raise a SCSI Unit Attention to make the host aware of the change. Workaround: No workaround is currently available.

Ticket ID	Severity	Affected versions	Description
SYS-266600	Moderate	11.5.0 or later	If a mirroring-related snapshot (LRS) of an active <code>xmirror</code> relation is deleted on the source volume due to lack of space in the storage pool, subsequent <code>xmirror_activate</code> commands on the source volume fail with a misleading return code (<code>XMIRROR_SLAVE_HAS_MORE_RECENT_DATA</code>). Workaround: 1. Remove the asynchronous source-destination mirror relation. 2. Ensure that the storage pool has enough space on the source. 3. Reestablish the asynchronous source-destination mirror relation. 4. Reestablish and activate the <code>xmirror</code> relation. During this process, the source-secondary synchronous mirror relation can continue working.
SYS-266620	Moderate	11.5.0 or later	The multi-tenancy feature does not support Sun directory or Open LDAP group objects. Static groups reflecting domain administrators fail to log in, and receive an unrecognized role completion code. Workaround: Instead of the an LDAP group object attribute, use the description string as hard-coded value per user. For more information, refer to the <i>IBM XIV Storage System Architecture and Implementation</i> Redbook.
SYS-267284	Moderate	11.6.0 or later	When a compressed volume is in a three-way mirroring relationship, any crash of the source volume might result in a situation where the tertiary site volume has more recent data than on the secondary site volume. In such a case, the regular recovery process does not work. Workaround: To resolve this issue: 1. Disable the three-way mirroring. 2. Recreate asynchronous mirroring with the tertiary site mirrored to the primary site. 3. Wait for the volume to synchronize. 4. Reestablish the three-way mirroring.
SYS-269298	Moderate	11.6.0 or later	The time required for decompressing a volume is proportional to the volume size, and the time required for compressing a volume is proportional to the amount of data that the volume contains. As a result, decompressing a volume is typically much slower than compressing the same volume. Workaround: No workaround is currently available.
SYS-257063	Service	11.3.0 or later	When using IBM Hyper-Scale Mobility (volume migration) and a host is mapped to both the source volume and the destination volume, the system does not display a warning upon the unmapping of the source volume from the host. Workaround: Prior to unmapping the source volume, verify that: • The host is mapped to the volume on the destination system. • I/Os are running on that path. This can be done from the host side or by using the XIV statistic tools on the destination system.
SYS-262628	Service	11.3.0 or later	Deleting a user certificate reverts the system to the generic certificate instead of to the unique system certificate. Workaround: Set the system-unique certificate as a default before removing the user certificate.
SYS-263091	Service	11.4.0 or later	When issuing the <code>encrypt_enable</code> XCLI command immediately after defining a key server: • The command returns <code>Command executed successfully</code>. • Encryption is not enabled. • The <code>ENCRYPT_ENABLE_NOT_COMPLETED</code> event is issued. Workaround: Wait 10 seconds and then retry.
SYS-263955	Service	11.5.0 or later	In a multi-site mirroring relation, if the secondary source is deleted by using the <code>delete_xmirror</code> command with <code>force=yes</code>, a subsequent <code>xmirror_activate</code> command might return a <code>success</code> return code, even though no secondary source volume is defined. Workaround: No workaround is currently available.
SYS-264016	Service	11.5.0 or later	Domain administrators from different domains that share a resource, can view and delete the metadata of that resource. Workaround: No workaround is currently available.
SYS-264238	Service	11.3.0 or later	<code>snmpwalk</code> for volume statistics may take a long time when multiple volumes are defined. Workaround: No workaround is currently available.
SYS-264393	Service	11.5.0 or later	When associating separate LDAP groups to separate domains, an LDAP directory user cannot belong to more than one group if not mentioned specifically by the group name in the format of <code>user@domain1/domain2</code>. If a user is not defined in such a way and the <code>memberOf</code> field contains more than one group, an <code>LDAP_USER_CONFIG_ERROR</code> event is issued. Workaround: No workaround is currently available.
SYS-264934	Service	11.5.0 or later	In IP filtering, FQDN host names are not supported by the <code>ip_access_group_remove_address</code> command. Workaround: Use an IP address instead.
SYS-265116	Service	11.5.0 or later	When using the multi-tenancy feature, the system does not block application administrators, who belong to user groups in the global domain, from being designated as IBM Hyper-Scale Manager users. Workaround: Remove any application administrator from the list of IBM Hyper-Scale Manager users.

Ticket ID	Severity	Affected versions	Description
SYS-265321	Service	11.5.0 or later	When working with IP filtering, a subnet cannot be added to an IP access group if an IPO, that is part of the subnet, already exists in the access group. Although the <code>ip_access_group_add_address</code> command is issued without an error, the subnet is not added. Workaround: No workaround is currently available.
SYS-265508	Service	11.5.0 or later	CIM is not multi-tenant aware and cannot be used by a domain administrator. Workaround: Use the credentials of a global (non-domain) user, and set the domain access policy to <code>OPEN</code> (default policy).
SYS-265642	Service	11.5.0 or later	IP filtering does not work for an IP range. Workaround: Define an IP range by using an IP mask.
SYS-265833	Service	11.5.0 or later	When using IP filtering, there is no validation that a valid IP address is entered to the IP Access Group. Workaround: No workaround is currently available.
SYS-265834	Service	11.5.0 or later	When adding IP addresses to an IP Access Group for IP filtering, two addresses cannot be added with the same IP, but with different netmasks. Workaround: No workaround is currently available.
SYS-269313	Service	11.5.0 or later	When three-way mirroring is being used, trying create a mirror snapshot on the asynchronous secondary system results in a failure (error event: <code>XMIRROR_SNAPSHOT_MIRROR_NOT_SUPPORTED</code>) because such operation is not supported with three-way mirroring. The prompted error message, however, is incorrect: Volume is part of an <code>xmirror</code> , snapshot mirrors are not supported for <code>xmirrors</code> . Snapshot mirrors are supported on a synchronous secondary system. Workaround: No workaround is currently available.
SYS-271583	Service	11.6.0 or later	1024-bit SSH keys that were in use prior to a software upgrade to version 11.6.x are no longer valid because 2048-bit SSH keys are required starting from version 11.6.x. Trying to use an 1024-bit SSH results in the following error message: <code>WARNING: REMOTE HOST IDENTIFICATION HAS CHANGED!</code> Workaround: Remove the old SSH key from the <code>\$HOME/.ssh/known_hosts</code> file, or delete the file entirely.
SYS-266619	Service	11.6.0 or later	When decompressing a compressed volume that belongs to a performance class, the decompression might take longer time to complete because the volume performance is constrained by the performance class. Workaround: Remove the performance class assignment until the conversion completes.
SYS-267859	Service	11.6.0 or later	When a compressed VMFS5 volume is being converted back to uncompressed state, the volume does not support VMware VAAI SCSI commands. These commands become supported again after the conversion process is complete. Workaround: To avoid this issue, the VMFS5 volume type must first be converted from ATS Only to ATS + SCSI reservation. Use the following ESXi CLI command to disable ATS locking: <code>vmkfstools --configATSOnly Odevice-path</code> For more information, see VMware KB article 2006858 (kb.vmware.com/kb/2006858).
SYS-268156	Service	11.6.0 or later	Obtaining the initial compression analysis of potential capacity saving for 1000 volumes might take up to 48 hours, depending on the load in the system. Workaround: No workaround is currently available.
SYS-268244	Service	11.6.0 or later	When volume compression is enabled, the general expectation would be that the cache hit ratio should increase, since some of the data is compressed and therefore the cache can contain more data. However, depending on the workload mix, the cache hit ratio might start to decline after a certain system load is exceeded. Workaround: Reduce the system load. For example, migrate some volumes to less loaded systems.
SYS-268876	Service	11.6.0 or later	When changing or switching a mirror relation role while compressed volumes are used, after the secondary system becomes a the primary system, the correct compression rate for the new primary system is not shown immediately. Workaround: No workaround is currently available.
SYS-269020	Service	11.6.0 or later	Cache hit latency for compressed volumes, as displayed in XIV TOP, is consistently incorrect and much higher than the actual one. Workaround: In the XIV performance view, use an aggregate filter and do not select the Hit or Miss filters. For higher accuracy, use external tools to measure the latency on the host.
SYS-269340	Service	11.6.0 or later	In some rare cases, upon the enabling of drive encryption, a solid-state drive (SSD) might fail, resulting in <code>COMPONENT_TEST_OF_SSD_HAS_FAILED</code> event for this drive. The SSD cache on the related module remains unavailable until the problem is fixed. Workaround: No workaround is currently available. For a possible solution, contact IBM Support.
SYS-269259	Service	11.6.0 or later	In a Microsoft Azure Site Recovery configuration, when a Failover operation is followed by Disable Protection, Microsoft System Center Virtual Machine Manager (SCVMM) might prompt an exception such as: <code>Replication group '\$cgname' was unexpectedly found to be in a replication relationship</code> . The exception is the result of a failure to automatically delete the mirror relationship on the target location. This might happen even if the communication between the source and target sites is operational. Workaround: Using XIV CLI or GUI, delete the consistency group mirror from the storage array on the target location before invoking Disable Protection. This can be applied before or after the exception occurs.
SYS-274880	Service	11.3.0 or later	Error code <code>MCL_TIMEOUT_EXPIRED</code> is returned when attempting to use the <code>cg_move</code> command to move a consistency group that is part of a group pool. Workaround: No workaround is currently available.

Ticket ID	Severity	Affected versions	Description
SYS-315110	Low	11.5.0 or later	The ipinterface_run_arp and ipinterface_list_ips XCLI commands do not support iSCSI (IP-based) interface queries. Workaround: Use the ipinterface_list XCLI command for iSCSI interface queries.

Release notes for IBM Hyper-Scale Manager

For release information about IBM® Hyper-Scale Manager, see the IBM Hyper-Scale Manager documentation.

- [IBM Hyper-Scale Manager documentation](https://ibm.com/docs/en/hyper-scale-manager) (ibm.com/docs/en/hyper-scale-manager)

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](https://www.ibm.com/support/fixcentral) 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](https://ibm.com/support/pages/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](https://ibm.com/support/pages/node/707025) (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

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 XIV Management Tools

The following release information is available for version 4.8.0.6 of IBM® XIV® Management Tools.

Important: As of June 2022, the IBM XIV Management Tools for XIV Storage System is no longer being maintained. This includes the following : XIV Management Tools 4.8.0.6 and lower, XIV XCLI 4.8.0.6 and lower, and IBM Hyper-Scale Manager v1.9.6 and lower. If these utilities are currently being used, it is important to migrate to the latest IBM Hyper-Scale Manager and XCLI Utility for the most up-to-date security patches. These can both be downloaded from [Fix Central](#). For more information about the XCLI Utility, see [Release notes for IBM XCLI Utility](#) and [Command-line reference](#). For more information about Hyper-Scale Manager, see [Hyper-Scale Manager documentation](#).

What's new in 4.8.0.7

IBM XIV Management Tools 4.8.0.7 is a minor release that introduces new operating system support and resolves a common security issue.

For more information, see [Requirements](#) and [Change log](#).

Compatibility

This section details the compatibility of IBM® XIV® Management Tools 4.8.0.7 with storage system versions and IBM Hyper-Scale Manager releases.

Supported storage systems

Version 4.8.0.7 supports the following storage system versions:

Storage system	Microcode version
IBM XIV Storage System Gen3 (281x-214, and 281x-314)	11.6.2.x
IBM Spectrum Accelerate	11.5.3.c, 11.5.4

Supported IBM Hyper-Scale Manager releases

IBM XIV Management Tools 4.8.0.7 is compatible with IBM Hyper-Scale Manager 1.9.7.

Requirements

This section details the XIV® GUI and IBM® Hyper-Scale Manager requirements.

- [XIV GUI requirements](#)
The XIV GUI runs in either Direct mode or Server mode, and requires different hardware resources depending on the mode in use.
- [IBM Hyper-Scale Manager requirements](#)
IBM XIV Management Tools 4.8.0.7 is bundled with IBM Hyper-Scale Manager 1.9.7, which can be installed as a server application.

XIV GUI requirements

The XIV® GUI runs in either Direct mode or Server mode, and requires different hardware resources depending on the mode in use.

Additionally, the requirements vary depending on whether the XIV GUI that is installed on a 32-bit or a 64-bit operating system.

Direct mode

In Direct mode, the XIV GUI requires the following:

CPU

Dual core

RAM

Up to 25 storage systems

32-bit and 64-bit: 700 MB

More than 25 storage systems

32-bit: 1.1 GB

64-bit: 1.5 GB

Max number of systems that can be monitored by a single XIV GUI in Direct mode:

144

Best practice: Use IBM® Hyper-Scale Manager and an XIV GUI in Server mode when using more than 10 storage systems.

Server mode

In Server mode, the XIV GUI requires the following:

CPU

Dual core

RAM

1 GB

Allowed latency between a single XIV GUI and IBM Hyper-Scale Manager:

Maximum of 50 milliseconds

Supported operating systems

Note: Make sure you install the relevant bit package (64-bit or 32-bit) according to the operating system that is in use.

Operating system	Bit version / CPU architecture	
Microsoft Windows 10	32-bit or 64-bit	
Red Hat Enterprise Linux® (RHEL) 7.x, 8.x	32-bit or 64-bit	
Apple OS X 12.x	N/A	
Operating system	Bit version / CPU architecture	Supported IBM XIV Management Tools component
Microsoft Windows Server 2008	32-bit or 64-bit	• XIV GUI • XIV Top • XIV CLI (XCLI) utility
Microsoft Windows Server 2016	64-bit	• XIV GUI • XIV Top • XIV CLI (XCLI) utility
Microsoft Windows Server 2012 R1 or R2	64-bit	• XIV GUI • XIV Top • XIV CLI (XCLI) utility
Microsoft Windows Vista	32-bit or 64-bit	• XIV GUI • XIV Top • XIV CLI (XCLI) utility
Microsoft Windows 7	32-bit or 64-bit	• XIV GUI • XIV Top • XIV CLI (XCLI) utility

Operating system	Bit version / CPU architecture	Supported IBM XIV Management Tools component
Microsoft Windows 8–8.1	32-bit or 64-bit	- XIV GUI - XIV Top - XIV CLI (XCLI) utility
Microsoft Windows 10	32-bit or 64-bit	- XIV GUI - XIV Top - XIV CLI (XCLI) utility
Red Hat Enterprise Linux (RHEL) 6.x	32-bit or 64-bit	- XIV GUI - XIV Top - XIV CLI (XCLI) utility
Apple OS X 10.9 – 10.12.x	N/A	- XIV GUI - XIV Top - XIV CLI (XCLI) utility
HP Unix (HP-UX) 11i v3	Intel IA-64	XCLI utility
IBM AIX® 7.0 and 7.1	N/A	XCLI utility
Solaris 11	SPARC or Intel x86-64	XCLI utility

IBM Hyper-Scale Manager requirements

IBM® XIV® Management Tools 4.8.0.7 is bundled with IBM Hyper-Scale Manager 1.9.7, which can be installed as a server application.

Note:

- Starting from version 1.7, IBM Hyper-Scale Manager as a virtual appliance is no longer supported. Only the application deployment is supported.
- If you are using the a previous virtual appliance version, refer to the *IBM Hyper-Scale Manager as an Application User Guide (GC27-5984)* for information on how to keep using the virtual appliance or migrate its data to the application version.

IBM Hyper-Scale Manager 1.9.7 requires the following software, hardware, and configurations:

Operating system

- Red Hat Enterprise Linux® (RHEL) 6.x7.x - 8.x, 64-bit.

Disk space

- At least 1 GB of free space in the /tmp directory
- 76 GB free space under the installation path (the path is configurable and does not include the home directory)

Physical CPU cores and memory

Number of storage systems in use	Number of storage objects in use	Required CPU cores	Required RAM
Up to 40	Up to 80,000	Dual-core	4 GB
Up to 80	Up to 120,000	4 cores	6 GB
Up to 120	Up to 180,000	6 cores	6 GB
Up to 144	Up to 220,000	6 cores	8 GB

Note: For best performance, install the IBM Hyper-Scale on a dedicated independent server.

Ports

- Outbound/Inbound: RMI - 1199 - 1204
- Inbound only: SSH/SCP - 22
- Outbound only: SMTP - 25
- Communication with storage systems - 7778
- RESTful API port - 8443

Latency

Allowed latency between the IBM Hyper-Scale Manager and the storage systems it monitors: up to 200 milliseconds.

Concurrently connected GUI clients to an XIV system

Maximum number of GUI clients (connected to a storage system) that can concurrently send requests to IBM Hyper-Scale Manager: 15

Administered storage systems

Maximum storage systems that are administered by one server: 144 (up to 12 groups with 12 storage systems in each)

Hyper-Scale Manager with XIV GUI

Use of IBM Hyper-Scale Manager is recommended when using 10 or more storage systems.

IBM Spectrum Accelerate

IBM Hyper-Scale Manager is required for IBM Spectrum Accelerate license management.

Change log

This change log summarizes the changes that were made in recent 4.x code level versions of IBM® XIV® Management Tools.

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.

- [4.8.0.7 \(March 2022\)](#)
IBM XIV Management Tools 4.8.0.7 resolves the following issue.
- [4.8.0.6 \(July 2017\)](#)
Version 4.8.0.6 of IBM XIV Management Tools resolves a few miscellaneous issues, as listed in the following table.
- [4.8.0.4 \(December 2016\)](#)
Version 4.8.0.4 of IBM XIV Management Tools introduced support for a new system authentication code and resolved a few miscellaneous issues, as listed in the following table.
- [4.8.0.3 \(July 2016\)](#)
Version 4.8.0.3 of IBM XIV Management Tools introduced a number of security enhancements and resolved the following issues.
- [4.8.0.2 \(April 2016\)](#)
Version 4.8.0.2 of IBM XIV Management Tools resolved the following issues.
- [4.8.0.1 \(January 2016\)](#)
Version 4.8.0.1 of IBM XIV Management Tools resolved the following issues.
- [4.8 \(December 2015\)](#)
Version 4.8 of IBM XIV Management Tools added support for IBM Spectrum Accelerate version 11.5.3 and IBM XIV Storage System Gen3 version 11.6.1.

4.8.0.6 (July 2017)

Version 4.8.0.6 of IBM® XIV® Management Tools resolves a few miscellaneous issues, as listed in the following table.

Ticket ID	Description
UI-265855	Fixed: An error occurs when you compress a Volume with an estimated compression ratio of less than 10% on a system with microcode version 11.6.2.a.
UI-265850	Fixed: The XIV GUI doesn't display on 2560*1440 high resolution screens.
UI-265840	Fixed: Data migration to an uncompressed Pool, when automatic Volume creation is enabled, results in configuration errors.

4.8.0.4 (December 2016)

Version 4.8.0.4 of IBM® XIV® Management Tools introduced support for a new system authentication code and resolved a few miscellaneous issues, as listed in the following table.

Ticket ID	Description
UI-265832	Security enhancement: When the password or username are incorrect during login, the XCLI return code does not disclose which information is incorrect.
UI-265697	Fixed: When a user is associated to a domain, all of the system sessions are disconnected for a few seconds, even if they are logged in as Admin.

4.8.0.3 (July 2016)

Version 4.8.0.3 of IBM® XIV® Management Tools introduced a number of security enhancements and resolved the following issues.

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 4.8.0.3, 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.
UI-265813	Fixed: The correct snmp_type cannot be defined, if the GUI is set to Japanese language mode.
UI-265811	Fixed: The units of measurement for bandwidth limits and bandwidth usage in the Statistics view in the Qos Performance Class Properties window are incorrect.

4.8.0.2 (April 2016)

Version 4.8.0.2 of IBM® XIV® Management Tools resolved the following issues.

Ticket ID	Description
UI-265802	Fixed: IBM Spectrum Accelerate deployment now supports a 6-digit ICN. (In case your ICN is 6 digits long, add a leading 0 to the number.)
UI-265800	Fixed: Sometimes the XIV GUI can't monitor pre-11.3.1 XIV systems.
UI-265799	Fixed: Wrong error message is displayed when creating a compressed volume (on 314 models with version 11.6.1).

4.8.0.1 (January 2016)

Version 4.8.0.1 of IBM® XIV® Management Tools resolved the following issues.

Ticket ID	Description
UI-265781	Fixed: The CVE-2015-7575 issue, also known as the "MD5 SLOTH" vulnerability. Management tools applications (Hyper-Scale Manager, XIV GUI, XIV Top) disable the MD5 algorithm.
UI-265787	Fixed: Wrong latency values are shown in XIV Top when the host is using a multipath.
UI-265634	Fixed: Wrong IOPS and latency values are shown when the system has more than six interface modules.

4.8 (December 2015)

Version 4.8 of IBM® XIV® Management Tools added support for IBM Spectrum Accelerate version 11.5.3 and IBM XIV Storage System Gen3 version 11.6.1.

Fixes in version 4.8:

Ticket ID	Description
UI-265713	Fixed: Removed users are still able to create consistency groups.

Fixes and features of earlier versions

IBM® XIV® Management Tools 4.8.0.7 includes all the features and fixes that were included in previous 4.x versions.

To obtain information regarding previously introduced features or issues that were resolved in previous versions, see the [Release notes publication page](#) Release notes publication page for your storage system IBM Documentation (ibm.com/docs/en).

Limitations

The following functional limitations apply to IBM® XIV® Management Tools and may remain in effect in future software versions.

Ticket ID	Description
UI-263912	After upgrading IBM Hyper-Scale Manager, the time set in Schedule Daily Backup is not saved and reverts to the default value. The default value is 11:30 PM.
UI-263821	The XIV GUI does not launch when Single Sign-On (SSO) software is running on the client.
UI-262890	Domains with size 0 cannot be created or edited.
UI-261561	An untrusted system (a system that is not authenticated via a certificate) is displayed as trusted. Removing the system certificate with the xcli -c remove command and then trying to connect to this system succeeds. This occurs when the XCLI server is up. Workaround: Log out of the XCLI server and then log in again.
UI-253185	The GUI must be installed in a dedicated directory, as the uninstallation erases it completely.
UI-247041	Changing the PC local time when the GUI is open might also change the storage system time. Workaround: Restart the GUI.


Known issues

This section details the known issues in IBM® XIV® Management Tools, along with possible solutions or workarounds (if available).


Important:

- **The issues listed below apply to version 4.8.0.7 or earlier versions.** As long as a newer version has not yet been released, a newer release notes edition for version 4.8.0.7 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 4.8.0.7 will no longer be updated. Accordingly, check the release notes of the newer version to learn whether any newly discovered issues affect version 4.8.0.7, or whether the newer version resolves any of the issues listed below.

Ticket ID	Description

Ticket ID	Description
UI-265880	The uninstall process of the XIV GUI fails on Microsoft Windows Server 2016 and Windows 10 operating systems. Workaround: 1. Access the Advanced system settings from the Windows Control Panel > System menu. 2. From the Advanced tab, click Environment Variables. 3. Click New to define a New User Variable. 4. Enter <code>JAVA_TOOL_OPTIONS</code> for the Variable name: and <code>-Dos.name="Windows 7"</code> for the Variable value:.
UI-265868	When collecting X-Ray support logs in the XIV GUI in Japanese, the logs can't be saved locally. Workaround: Use the XIV GUI in English.
UI-265863	The installation process of Hyper-Scale Manager version 1.9.4 fails. Workaround: After the Hyper-Scale Manager installation package is extracted into the <code>tmp</code> folder, temporarily mount the folder without the <code>noexec</code> flag and install Hyper-Scale Manager.
UI-265848	The XIV GUI doesn't support SED on IBM Spectrum Accelerate systems. Workaround: No workaround is currently available.
UI-265847	In the XIV GUI Japanese interface, the Compression Capabilities field in the Parameters tab of the system settings is disabled. Workaround: Use the XCLI or change the interface to English.
UI-265846	An error is displayed when you automatically collect and upload support logs to ECUREP from the XIV GUI. Workaround: Manually upload the log files via FTP.
UI-265838	An error occurs when trying to launch the XIV GUI after an installation on OSX 10.12.x. Workaround: 1. Edit the <code>Info.plist</code> file in the XIV GUI.app/Contents directory. 2. Replace <code>%BUNDLE_PARENT_PATH%</code> with <code>/Applications/XIVGUI</code> throughout the file. For example, instead of <code><string>%BUNDLE_PARENT_PATH%/lib/</string></code>, you should have <code><string>/Applications/XIVGUI/lib/</string></code>. Note: Be sure to make the same updates for XCLI and XIV Top.
UI-265789	When trying to import a custom certificate, the custom certificate is overwritten with the original self-signed certificate. Workaround: Contact IBM Support.
UI-265788	After a successful upgrade of HSM from version 1.9 to 1.9.1, the systems might get stuck in the Fetching Data status. Workaround: Reopen the XIV GUI.
UI-265747	In the XIV GUI 4.8 build 11 installer on Mac OSX 10.11 (El Captain), the legends are missing on all four buttons in the License Agreement dialog. In addition, the red (Close) button in the dialog menu bar is missing.  The buttons (from left to right) are: Print, Save, Decline, Accept.

Ticket ID	Description
UI-265714	On the LDAP configuration panel of the management GUI, the Storage Integration Admin Role field is incorrectly marked as mandatory, and the LDAP configuration cannot be completed without entering a value in that field. The fact that this field is not mandatory is correctly documented in the <i>IBM XIV Storage System Command-line Interface (CLI) Reference Guide</i> and the <i>IBM Spectrum Accelerate Command-line Interface (CLI) Reference Guide</i>. Workaround: Enter a fictitious value in the Storage Integration Admin Role field and save the LDAP configuration. Then, remove the role value with the command-line interface (CLI) by using the following command: <pre>ldap_config_set storage_integration_admin_role=""</pre>
UI-265379	An asynchronous mirror relation with an RPO that is different from the default (30 seconds) cannot be extended to a three-way mirroring with the standby mirror option selected. In addition, an existing synchronous mirror relation cannot be extended to three-way mirroring if the asynchronous mirror created during the process has an RPO that is different from the default. Workaround: Manually create the mirror schedule on the secondary source and destination systems (if a standby mirror is created). Run the scheduler_create CLI command prior to creating the three-way mirroring. Alternatively, you can set the RPO with the default value and modify it after the three-way mirroring is successfully created.
UI-265088	On the target of a compressed volume, that is either mirrored or part of a Hyper-Scale Mobility (volume migration) relationship, the ratio and savings parameters are wrongly displayed as zero (0). Workaround: No workaround is currently available.
UI-264605	When there is a role conflict in a mirroring relation, two rows are displayed in the GUI. The rows display the storage systems with their source and destination volumes in different (and opposite) roles. That is, the system's role in the mirroring relation is reversed in each row. Workaround: Select only the row on which you would like to fix the role conflict and run the Change Role option on the mirror's source and destination. Refrain from selecting both rows at once, as this will only lead to further problems.
UI-264145	The XIV GUI truncates (rounding down) the Latency values in the Statistics view. Workaround: No workaround is currently available.
UI-264141	When there are two identical rows for the same mirror, the cursor might jump to the first mirror when you try to select the bottom mirror. Workaround: No workaround is currently available.
UI-263084	When an XIV Top session is open and the session idle timeout is changed, it will only take effect the next time the XIV Top is opened. Workaround: Close and reopen a new XIV Top session.
UI-261999	The option to determine where an imported certificate is stored does not work. The certificate is stored in C:\Users\{UserName}\AppData\Roaming\XIV\GUI12\properties, regardless of the address that is stated by the -h parameter. Workaround: Import the certificate using IBM Hyper-Scale Manager.
UI-243418	The XIV GUI does not launch if certain env variables are defined. The specific env variables are: • IBM_JAVA_OPTIONS • JAVA_TOOL_OPTIONS Workaround: Delete the definition.
MSM-261090	IBM Hyper-Scale Manager cannot monitor pre-11.2.0 XIV systems that are configured with LDAP with an <i>xiv_msms</i> user that is configured both internally and in the LDAP. Workaround: 1. Disable LDAP on the specific XIV system. <pre>ldap_mode_set mode=inactive</pre> 2. Delete the <i>xiv_msms</i> user which is defined internally. <pre>user_delete user=xiv_msms</pre> 3. Re-enable LDAP. <pre>ldap_mode_set mode=active</pre>

Ticket ID	Description
MSM-261538	Disabling LDAP on an XIV system may cause an authentication failure on all of the systems that are configured to this LDAP. Workaround: 1. Select Systems > Manager Configuration from the XIV GUI menu. 2. Type the Manager Access Code. 3. Select the Inventory tab. 4. Select and right-click the systems whose status is Authentication Failure and click Diagnose/Fix Authentication Error on the menu.  The screenshot shows the 'Manager Configuration' window with the 'Inventory' tab selected. A table lists 15 systems. The system 'XIV Gen3P2-71' is highlighted in yellow and has a status of 'Authentication Failure'. A right-click context menu is open over this system, displaying the following options: Add System, Add Group, Move System To ..., Delete, Edit, Suspend Monitoring, Enable Monitoring, Up, Down, Diagnose/Fix Authentication Error, and Show Certificate. A red text overlay at the bottom of the window reads: 'Some of the systems are not authenticated. Right click then'.

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](https://ibm.com/support/pages/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](https://ibm.com/support/pages/node/707025) (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 <i>noexec</i> 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 <i>exec</i> option, then remount it with <i>noexec</i> 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 notes PDF publications

Use this page to access PDF versions of release notes for IBM XIV Storage System Gen3, IBM XIV Management Tools, IBM Extended Command-Line Interface (XCLI) Utility, and IBM Remote Support Proxy.

- [IBM XIV Storage System release notes](#)
- [IBM XIV Management Tools 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](#).

Note: For Hyper-Scale Manager, including REST API, publications, see [IBM Hyper-Scale Manager documentation](#).

IBM XIV Storage System Gen3 release notes

Release	Edition	PDF download link
11.6.2.d	First edition (August 2019)	 Download
11.6.2.c	First Edition (August 2018)	 Download
11.6.2.b	First Edition (May 2018)	 Download
11.6.2.a	Third Edition (February 2018)	 Download
11.6.2	Second Edition (January 2017)	 Download
11.6.1.a	Third Edition (September 2016)	 Download
11.6.1	Third Edition (January 2016)	 Download
11.6.0.b	Fifth Edition (January 2016)	 Download
11.5.2.b	First Edition (May 2016)	 Download
11.5.2.a	First Edition (March 2016)	 Download
11.5.2	First Edition (August 2015)	 Download

IBM XIV Management Tools release notes

Important: As of June 2022, the IBM XIV Management Tools for XIV Storage System is no longer being maintained. This includes the following : XIV Management Tools 4.8.0.6 and lower, XIV XCLI 4.8.0.6 and lower, and IBM Hyper-Scale Manager v1.9.6 and lower. If these utilities are currently being used, it is important to migrate to the latest IBM Hyper-Scale Manager and XCLI Utility for the most up-to-date security patches. These can both be downloaded from [Fix Central](#). For more information about the XCLI Utility, see [Release notes for IBM XCLI Utility](#) and [Command-line reference](#). For more information about Hyper-Scale Manager, see [Hyper-Scale Manager documentation](#).

Release	Edition	PDF download link
4.8.0.6	Second Edition (July 2017)	 Download
4.8.0.4	Second Edition (March 2017)	 Download
4.8.0.3	First Edition (July 2016)	 Download
4.8.0.2	First Edition (April 2016)	 Download
4.8.0.1	Third Edition (March 2016)	 Download
4.8	Second Edition (December 2015)	 Download
4.7	First Edition (September 2015)	 Download
4.6.0.1	First Edition (June 2015)	 Download
4.6	First Edition (June 2015)	 Download
4.5.0.2	First Edition (June 2015)	 Download
4.5.0.1	First Edition (April 2015)	 Download
4.5	First Edition (March 2015)	 Download
4.4.1.1	First Edition (February 2015)	 Download
HSM 1.5.1.33	First Edition (February 2015)	 Download
HSM 1.5.1	December 2014	 Download
4.4.1	December 2014	 Download
4.4.0.2	October 2014	 Download
HSM 1.5.0.59	October 2014	 Download
4.4	August 2014	 Download
HSM 1.5	July 2014	 Download
4.3.1	January 2014	 Download
4.3	Second Edition (November 2013)	 Download
HSM 1.4	November 2013	 Download
4.2.2	July 2013	 Download
4.2.1	July 2013	 Download
4.2	June 2013	 Download
MSM 1.3	June 2013	 Download
MSM 1.2	March 2013	 Download
4.1	March 2013	 Download
MSM 1.1	November 2012	 Download

Release	Edition	PDF download link
4.0	November 2012	 Download

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
5.1	Second Edition (March 2017)	 Download
5.0	Second Edition (July 2016)	 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

IBM XIV Storage System overview

The IBM® XIV® Storage System is a next-generation high-end open disk storage system.

An integral part of IBM's broad spectrum of system storage offerings, the XIV system has an innovative grid architecture designed to deliver the highest levels of reliability, performance, scalability and functionality at low overall cost, while eliminating complexity and providing unprecedented ease of management.

Designed for exceptionally low total cost of ownership, the XIV system aims to tackle almost every aspect related to owning, operating and maintaining a storage system, including backup and restore, acquisition, environmental, administration, and downtime cost factors.

The XIV system is designed to deliver consistently high performance through the elimination of hot spots and full exploitation of all key system resources at all times. Its ground-breaking architecture delivers these performance levels while innovatively leveraging the capacity and cost benefits of Very High Density Slower Rotation (VHDSR) drives.

Outstanding availability and reliability are delivered through a revolutionary redundancy scheme, enabling automated self healing with exceptionally fast rebuild times. Software functionality including snapshotting, thin provisioning and space reclamation, and remote mirroring can help ease provisioning and protect business data from unexpected outages. IBM XIV can help customers deploy an information infrastructure that is reliable, versatile and available while offering improvements in storage management, provisioning and improved asset utilization.

This information can help you configure your storage, attach hosts, set up remote mirroring, and use the other features and functions of your IBM XIV Storage System. You can also find more information about the IBM XIV Storage System product at [IBM XIV Storage System: Storage Reinvented](#).

- [Introduction](#)
IBM XIV is a high-end grid-scale storage system that delivers consistently high performance, high resiliency and management simplicity while offering exceptional data economics, including powerful real-time compression. Industry benchmarks underscore stellar XIV performance and cost benefits.
- [Volumes and snapshots](#)
This section gives an overview of volumes and snapshots.
- [Storage pools](#)
The storage space of the IBM XIV Storage System is portioned into *storage pools*, where each volume belongs to a specific storage pool.
- [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.
- [QoS performance classes](#)
The Quality of Service (QoS) feature allows the IBM XIV Storage System to deliver different service levels to hosts that are connected to the same XIV system.
- [Connectivity with hosts](#)
The storage system connectivity is provided through the following interfaces:
- [IBM Real-time Compression with XIV](#)
This section introduces the IBM® Real-time Compression™ (RtC) feature of IBM XIV storage systems.
- [Synchronous remote mirroring](#)
IBM XIV Storage System features synchronous and asynchronous remote mirroring for disaster recovery. Remote mirroring can be used to replicate the data between two geographically remote sites. The replication ensures uninterrupted business operation if there is a total site failure.
- [Asynchronous remote mirroring](#)
Asynchronous mirroring enables you to attain high availability of critical data through a process that asynchronously replicates data updates that are recorded on a primary storage peer to a remote, secondary peer.
- [Multi-site mirroring](#)
Multi-site mirroring is an IBM XIV Storage System technology that allows customers to set High Availability and Disaster Recovery solutions over multiple sites, keeping 3 copies of their data.
- [IBM Hyper-Scale Mobility](#)
IBM Hyper-Scale Mobility enables a non-disruptive migration of volumes from one storage system to another.

- [Data-at-rest encryption](#)
The IBM XIV Storage System utilizes full disk encryption for regulation compliance and security audit readiness.
- [Management and monitoring](#)
The storage system can be monitored and fully controlled by using different management and automation tools.
- [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. Each of the following destinations must be defined as described:
- [User roles and permissions](#)
User roles allow specifying which roles are applied and the various applicable limits.
- [User authentication and access control](#)
IBM XIV Storage System features role-based authentication either natively or by using LDAP-based authentication.
- [Multi-Tenancy](#)
The storage system allows allocating storage resources to several independent administrators, assuring that one administrator cannot access resources associated with another administrator.
- [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.
- [Software upgrade](#)
Non disruptive code load (hot upgrade) enables the IBM XIV Storage System to upgrade its software from a current version to a newer version without disrupting application service.
- [Remote support and proactive support](#)
To allow IBM to provide support for the storage system, the proactive support and remote support options are available.

Introduction

IBM® XIV® is a high-end grid-scale storage system that delivers consistently high performance, high resiliency and management simplicity while offering exceptional data economics, including powerful real-time compression. Industry benchmarks underscore stellar XIV performance and cost benefits.

Figure 1. IBM XIV Storage System unit



As a grid-scale offering, every IBM XIV storage system contains multiple modules that are interconnected by integrated InfiniBand switches, forming a scale-out grid fabric that delivers exceptional IOPS performance. In addition, it includes a maintenance module for remote access to the system and an uninterruptible power supply modules to ensure system operation if an external power source fails.

IBM XIV storage system is ideal for cloud environments, offering predictable service levels for dynamic workloads, simplified scale management including in multi-tenant environments, flexible consumption models — and robust cloud automation and orchestration through OpenStack, RESTful API, and VMware. It offers data-at-rest encryption, advanced mirroring and self-healing, and provides investment protection with perpetual licensing that is transferable to all Spectrum Accelerate Family offerings (XIV Gen3, FlashSystem A9000/A9000R and Spectrum Accelerate).

- [Features and functionality](#)
IBM XIV Storage System is characterized by powerful features and functionality.
- [Hardware overview](#)
This section provides a general overview of the IBM XIV Storage System hardware.

- [Reliability](#)
IBM XIV Storage System reliability features include data mirroring, spare storage capacity, self-healing mechanisms, and data virtualization.
- [Performance](#)
The IBM XIV Storage System is a ground breaking, high performance storage product designed to help enterprises overcome this challenge through an exceptional mix of game-changing characteristics and capabilities
- [Functionality](#)
IBM XIV Storage System functions include point-in-time copying, automatic notifications, and ease of management through a GUI or CLI.

Features and functionality

IBM® XIV® Storage System is characterized by powerful features and functionality.

These features and functionality include:

Performance

- Perfect load balancing
- Cache and disks in every module
- Extremely fast rebuild time in the event of disk failure
- Constant, predictable high performance with zero tuning

Reliability

- Unique data distribution method that eliminates "hot spots"
- Fault tolerance, failure analysis, and self-healing algorithms
- No single-point-of-failure

Scalability

- Support for thin provisioning
- Support for instant space reclamation
- Data migration

Connectivity

- iSCSI and Fibre Channel (FC) interfaces
- Multiple host access

Snapshots

- Innovative snapshot functionality, including support for practically unlimited number of snapshots, snap-of-snap and restore-from-snap

Replication

- Synchronous and asynchronous replication of a volume (as well as a consistency group) to a remote system

Ease of management

- Support for storage pools administrative units
- Remote configuration management
- Non-disruptive maintenance and upgrades
- Management software, including a graphical user interface (GUI) and a command-line interface (CLI)
- Notifications of events through e-mail, SNMP, or SMS messages
- XIV is supported by the following IBM products:
 - IBM Power Virtualization Center (PowerVC). This is an advanced virtualization management offering that simplifies creating and managing virtual machines on IBM Power Systems™ servers using PowerVM® or PowerKVM hypervisors.
 - IBM Spectrum Protect Snapshot. Formerly Tivoli® Storage FlashCopy® Manager, IBM® Spectrum Protect™ Snapshot delivers high levels of protection for key applications and databases using advanced integrated application snapshot backup and restore capabilities.
 - IBM Spectrum Control. Formerly IBM Tivoli Storage Productivity Center, IBM Spectrum Control is integrated data and storage management software that provides monitoring, automation and analytics for organizations with multiple storage systems.


Hardware overview

This section provides a general overview of the IBM® XIV® Storage System hardware.

- [Hardware components](#)
The IBM XIV Storage System configuration includes data modules, interface modules, Ethernet switches, and uninterruptible power supply units.
- [Hardware enhancements](#)
IBM XIV Storage System Gen3 model 314 is a hardware-enhanced XIV Gen3 storage array targeted to customers who want high utilization of IBM Real-time Compression (RtC).
- [Management options](#)
The IBM XIV Storage System provides several management options.

Hardware components

The IBM® XIV® Storage System configuration includes data modules, interface modules, Ethernet switches, and uninterruptible power supply units.



Data modules

Each data module contains 12 disks DDR3 and 24GB of cache memory. IBM XIV supports all sorts of Near-line (7200RPM) SAS drives, in particular 2TB, 4TB, or 6TB disks. The disk drives serve as the nonvolatile memory for storing data in the storage grid and the cache memory is used for caching data previously read, prefetching of data from a disk, and for delayed destaging of previously written data.

Note: Data modules are located on both upper and lower sections of the rack.

InfiniBand switch (x2)

- Dual power supply
- Each IB switch is connected to each XIV module
- Both IB switches are inter-connected on their ports 16, 17
- Port 18 is left empty for spare
- Ports 19-36 are inactive

Maintenance module

Allows remote support access using a modem.

Interface modules (IM)

Each contains disk drives and cache memory similar to the data modules. In addition, these modules have Host Interface Adapters with FC and iSCSI ports. 8Gbps FC

- 2 dual-port FC HBAs on each interface module, i.e. 4 ports on each interface module, i.e. 24 ports on a full rack.

Uninterruptible power supply module complex

The uninterruptible power supply module complex consists of three units. It maintains an internal power supply in the event of a temporary failure of the external power supply. In the case of a continuous external power failure, the uninterruptible power supply module complex maintains power long enough for a safe and ordered shutdown of the IBM XIV Storage System. The IBM XIV Storage System can sustain the failure of one uninterruptible power supply unit while protecting against external power failures.

ATS

The Automatic Transfer Switch (ATS) switches between line cords in order to allow redundancy of external power.

Modem

Allows the system to receive a connection for remote access by IBM support. The modem connects to the maintenance module.

Data and interface modules are generically referred to as "modules". Modules communicate with each other by means of the PCIe adapter. Each module contains redundant ports for module to module communication. The ports are all linked to the internal network through the switches. In addition, for monitoring purposes, the UPSs are directly connected to the individual modules.

Hardware enhancements

IBM® XIV® Storage System Gen3 model 314 is a hardware-enhanced XIV Gen3 storage array targeted to customers who want high utilization of IBM Real-time Compression (RtC).

With double the RAM and CPU resources, IBM XIV Storage System model 314 delivers improved IOPS per compressed capacity and 1 to 2 PB of effective capacity without performance degradation. IBM XIV Storage System model 314 hardware enhancements include:

- 2 x 6-core CPUs per module (versus 1 x 6-core CPU per module in Model 214)
- 96 GB RAM per module (versus 48 GB RAM per module in Model 214)

Note: The additional CPU and 48 GB RAM are dedicated to Real-time Compression. For more information on Real-time Compression, see [IBM Real-time Compression with XIV](#).

Available configurations

IBM XIV Storage System model 314 is available for ordering in the following configurations:

- 9 to 15 modules in a system
- 4 TB or 6 TB drives
- 800 GB SSD cache (mandatory)

IBM XIV Storage System version 11.6.2 with IBM XIV Storage Gen3 System Model 314 also supports the following:

- Up to 2 PB of available soft capacity
- Reduced minimum compressible volume size from 103 GB (in model 214) to 51 GB
- Support for IBM Spectrum Accelerate software licenses

To learn more about IBM Real-time Compression, go to [IBM Real-time Compression with XIV](#).

For more information, see the *IBM XIV Storage System Release Notes*, version 11.6.1 documentation.

Management options

The IBM® XIV® Storage System provides several management options.

GUI and CLI management applications

These applications must be installed on each workstation that will be used for managing and controlling the system. All configurations and monitoring aspects of the system can be controlled through the GUI or the CLI.

SNMP

Third-party SNMP-based monitoring tools are supported using the IBM XIV MIB.

E-mail notifications

The IBM XIV Storage System can notify users, applications or both through e-mail messages regarding failures, configuration changes, and other important information.

SMS notifications

Users can be notified through SMS of any system event.

Reliability

IBM® XIV® Storage System reliability features include data mirroring, spare storage capacity, self-healing mechanisms, and data virtualization.

- [Redundant components](#)
IBM XIV Storage System hardware components are fully redundant, and ensure failover protection for each other to prevent a single point of system failure.
- [Data mirroring](#)
Data arriving from the host for storage is temporarily placed in two separate caches before it is permanently written to two disk drives located in separate modules. This guarantees that the data is always protected against possible failure of individual modules, and this protection is in effect even before data has been written to the nonvolatile disk media.
- [Self-healing mechanisms](#)
The IBM XIV Storage System includes built-in mechanisms for self-healing to take care of individual component malfunctions and to automatically restore full data redundancy in the system within minutes.
- [Protected cache](#)
- [Redundant power](#)
- [SSD cache drives](#)
The IBM XIV Storage System uses Flash-as-Cache (SSD) as a second, read-only, caching layer between the cache node and the disks.

Redundant components

IBM® XIV® Storage System hardware components are fully redundant, and ensure failover protection for each other to prevent a single point of system failure.

System failover processes are transparent to the user because they are swiftly and seamlessly completed.

Data mirroring

Data arriving from the host for storage is temporarily placed in two separate caches before it is permanently written to two disk drives located in separate modules. This guarantees that the data is always protected against possible failure of individual modules, and this protection is in effect even before data has been written to the nonvolatile disk media.

Self-healing mechanisms

The IBM® XIV® Storage System includes built-in mechanisms for self-healing to take care of individual component malfunctions and to automatically restore full data redundancy in the system within minutes.

Self-healing mechanisms dramatically increase the level of reliability in the IBM XIV Storage System. Rather than necessitating a technician's on-site intervention in the case of an individual component malfunction to prevent a possible malfunction of a second component, the automatically restored redundancy allows a relaxed maintenance policy based on a pre-established routine schedule.

Self-healing mechanisms are not just started in a reactive fashion following an individual component malfunction, but also proactively - upon detection of conditions indicating potential imminent failure of a component. Often, potential problems are identified well before they might occur with the help of advanced algorithms of preventive self-analysis that are continually running in the background. In all cases, self-healing mechanisms implemented in the IBM XIV Storage System identify all data portions in the system for which a second copy has been corrupted or is in danger of being corrupted. The IBM XIV Storage System creates a secure second copy out of the existing copy, and it stores it in the most appropriate part of the system. Taking advantage of the full data virtualization, and based on the data distribution schemes implemented in the IBM XIV Storage System, such processes are completed with minimal data migration.

As with all other processes in the system, the self-healing mechanisms are completely transparent to the user, and the regular activity of responding to I/O data requests is thoroughly maintained with no degradation to system performance. Performance, load balance, and reliability are never compromised by this activity.

Protected cache

IBM® XIV® Storage System cache writes are protected. Cache memory on a module is protected with ECC (Error Correction Coding). All write requests are written to two separate cache modules before the host is acknowledged. The data is later destaged to disks.

Redundant power

Redundancy of power is maintained in the IBM® XIV® Storage System through the following means:

- Three uninterruptible power supply units - the system can run indefinitely on two uninterruptible power supply units. No system component will lose power if a single uninterruptible power supply unit fails.
- Redundant power supplies in each data and interface module. There are two power supplies for each module and each power supply for a module is powered by a different uninterruptible power supply unit.
- Redundant power for Ethernet switches - each Ethernet switch is powered by two uninterruptible power supply units. One is a direct connect; one is through the Ethernet switch redundant power supply.
- Redundant line cords - to protect against the loss of utility power, two line cords are supplied to the ATS. If utility power is lost on one line cord, the ATS automatically switches to the other line cord, without impacting the system.
- In the event of loss of utility power on both line cords, the uninterruptible power supply units will maintain power to the system until an emergency destage of all data in the system can be performed. Once the emergency destage has completed, the system will perform a controlled power down.

SSD cache drives

The IBM XIV Storage System uses Flash-as-Cache (SSD) as a second, read-only, caching layer between the cache node and the disks.

This way, the system reduces disk access with having a cache that is an order-of-magnitude larger than the DRAM cache. Currently, the system features one SSD disk per module. Flash-as-cache is designed in module granularity, so lacking this feature in one module does not affect its functionality on other modules. Flash-as-cache can be enabled and disabled at run time, so a storage system can be equipped with SSDs anytime.

Installation

The SSD is a new hardware component that is identified as **1:SSD:<module:1>**. It is automatically added when plugged in (no equip command needed). The SSD can be phased-out, tested and phased-in like a regular disk drive. Since it is read-cache only, no rebuild is triggered due to SSD component state change. It is implemented and accessed similarly to other disks.

Performance

The IBM® XIV® Storage System is a ground breaking, high performance storage product designed to help enterprises overcome this challenge through an exceptional mix of game-changing characteristics and capabilities

Breakthrough architecture and design

The revolutionary design of IBM XIV Storage System enables exceptional performance optimization typically unattainable by traditional architectures. This optimization results in superior utilization of system resources and automatic workload distribution across all system hard drives. It also empowers administrators to tap into the system's rich set of built-in, advanced functionality such as thin provisioning, mirroring and snapshots without adversely affecting performance.

Consistent, predictable performance and scalability

The IBM XIV Storage System's ability to optimize load distribution across all disks for all workloads, coupled with a powerful distributed cache implementation, facilitates high performance that scales linearly with added storage enclosures. Because this high performance is consistent—without the need for manual tuning—users can enjoy the same high performance during the typical peaks and troughs associated with volume and snapshot usage patterns, even after a component failure.

Resilience and self-healing

The IBM XIV Storage System maintains resilience during hardware failures, continuing to function with minimal performance impact. Additionally, the solution's advanced self-healing capabilities allow it to withstand additional hardware failures once it recovers from the initial failure.

Automatic optimization and management

Unlike traditional storage solutions, the IBM XIV Storage System automatically optimizes data distribution through hardware configuration changes such as component additions, replacements or failure. This helps eliminate the need for manual tuning or optimization.

Functionality

IBM® XIV® Storage System functions include point-in-time copying, automatic notifications, and ease of management through a GUI or CLI.

- [Snapshot management](#)
The IBM XIV Storage System provides powerful snapshot mechanisms for creating point-in-time copies of volumes.
- [Consistency groups for snapshots](#)
Volumes can be put in a consistency group to facilitate the creation of consistent point-in-time snapshots of all the volumes in a single operation.
- [Storage pools](#)
Storage pools are used to administer the storage resources of volumes and snapshots.
- [Remote monitoring and diagnostics](#)
IBM XIV Storage System can email important system events to IBM Support.
- [SNMP](#)
Third-party SNMP-based monitoring tools are supported for the IBM XIV Storage System MIB.
- [Multipathing](#)
The parallel design underlying the activity of the Host Interface modules and the full data virtualization achieved in the system implement thorough multipathing access algorithms.
- [Automatic event notifications](#)
The system can be set to automatically transmit appropriate alarm notification messages through SNMP traps, or email messages.
- [Management through GUI and CLI](#)
The IBM XIV Storage System offers a user-friendly and intuitive GUI application and CLI commands to configure and monitor the system.
- [External replication mechanisms](#)
External replication and mirroring mechanisms in the IBM XIV Storage System are an extension of the *internal replication mechanisms* and of the overall functionality of the system.
- [Upgradability](#)
The IBM XIV Storage System is available in a partial rack system comprised of as few as six (6) modules, or as many as fifteen (15) modules per rack.

Snapshot management

The IBM® XIV® Storage System provides powerful snapshot mechanisms for creating point-in-time copies of volumes.

The snapshot mechanisms include the following features:

- Differential snapshots, where only the data that differs between the source volume and its snapshot consumes storage space
- Instant creation of a snapshot without any interruption of the application, making the snapshot available immediately
- Writable snapshots, which can be used for a testing environment; storage space is only required for actual data changes
- Snapshot of a writable snapshot can be taken
- High performance that is independent of the number of snapshots or volume size
- The ability to restore from snapshot to volume or snapshot

Consistency groups for snapshots

Volumes can be put in a consistency group to facilitate the creation of consistent point-in-time snapshots of all the volumes in a single operation.

This is essential for applications that use several volumes concurrently and need a consistent snapshot of all these volumes at the same point in time.

Storage pools

Storage pools are used to administer the storage resources of volumes and snapshots.

The storage space of the IBM® XIV® Storage System can be administratively portioned into storage pools to enable the control of storage space consumption for specific applications or departments.

Remote monitoring and diagnostics

IBM® XIV® Storage System can email important system events to IBM Support.

This allows IBM to immediately detect hardware failures warranting immediate attention and react swiftly (for example, dispatch service personnel). Additionally, IBM support personnel can conduct remote support and generate diagnostics for both maintenance and support purposes. All remote support is subject to customer permission and remote support sessions are protected with a challenge response security mechanism.

SNMP

Third-party SNMP-based monitoring tools are supported for the IBM® XIV® Storage System MIB.

Multipathing

The parallel design underlying the activity of the Host Interface modules and the full data virtualization achieved in the system implement thorough multipathing access algorithms.

Thus, as the host connects to the system through several independent ports, each volume can be accessed directly through any of the Host Interface modules, and no interaction has to be established across the various modules of the Host Interface array.

Automatic event notifications

The system can be set to automatically transmit appropriate alarm notification messages through SNMP traps, or email messages.

The user can configure various triggers for sending events and various destinations depending on the type and severity of the event. The system can also be configured to send notifications until a user acknowledges their receipt.

Management through GUI and CLI

The IBM® XIV® Storage System offers a user-friendly and intuitive GUI application and CLI commands to configure and monitor the system.

These feature comprehensive system management functionality, encompassing hosts, volumes, consistency groups, storage pools, snapshots, mirroring relationships, data migration, events, and more.

For more information, see the *IBM XIV Management Tools Operations Guide* and *IBM XIV Storage System XCLI User Manual*.

External replication mechanisms

External replication and mirroring mechanisms in the IBM® XIV® Storage System are an extension of the *internal replication mechanisms* and of the overall functionality of the system.

These features provide protection against a site disaster to ensure production continues. The mirroring can be performed over either Fibre Channel or iSCSI, and the host-to-storage protocol is independent of the mirroring protocol.

Upgradability

The IBM® XIV® Storage System is available in a partial rack system comprised of as few as six (6) modules, or as many as fifteen (15) modules per rack.

Partial rack systems may be upgraded by adding data and interface modules, up to the maximum of fifteen (15) modules per rack.

The system supports a non-disruptive upgrade of the system, as well as hotfix updates.

Volumes and snapshots

This section gives an overview of volumes and snapshots.

Volumes are the basic storage data units in the IBM® XIV® 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 and storage pools.

The basic hierarchy may be described 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.

The following subsections deal with volumes and snapshots specifically.

- [Volume function and lifecycle](#)
The volume is the basic data container that is presented to the hosts as a logical disk.
- [Snapshot function and lifecycle](#)
The roles of the snapshot determine its life cycle.
- [Additional snapshot attributes](#)
Snapshots have the following additional attributes.
- [Redirect-on-Write \(ROW\)](#)
The IBM XIV 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 master volume from any of its locked snapshots.

Volume function and lifecycle

The 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 *master 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 IBM® XIV® 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.

- [Support for Symantec Storage Foundation Thin Reclamation](#)
The IBM XIV Storage System supports Symantec's Storage Foundation Thin Reclamation API.

Support for Symantec Storage Foundation Thin Reclamation

The IBM® XIV® Storage System supports Symantec's Storage Foundation Thin Reclamation API.

The IBM XIV Storage System features instant space reclamation functionality, enhancing the existing IBM XIV Thin Provisioning capability. The new instant space reclamation function allows IBM XIV users to optimize capacity utilization, thus saving costs, by allowing supporting applications, to instantly regain unused file system space in thin-provisioned XIV volumes instantly.

The IBM XIV Storage System is one of the first high-end storage systems to offer instant space reclamation. The new, instant capability enables third party products vendors, such as Symantec Thin Reclamation, to interlock with the The IBM XIV Storage System such that any unused space is detected instantly and automatically, and immediately reassigned to the general storage pool for reuse.

This enables integration with thin-provisioning-aware Veritas File System (VxFS) by Symantec, which ultimately enables to leverage the IBM XIV Storage System thin-provisioning-awareness to attain higher savings in storage utilization.

For example, when data is deleted by the user, the system administrator can initiate a reclamation process in which the IBM XIV Storage System frees the non-utilized blocks and where these blocks are reclaimed by the available pool of storage.

Instant space reclamation does not support space reclamation for the following objects:

- Mirrored volumes
- Volumes that have snapshots
- Snapshots

Snapshot function and lifecycle

The roles of the snapshot determine its life cycle.


The IBM® XIV® 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 master 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.

With the IBM XIV snapshots, no storage capacity is consumed by the snapshot until the source volume (or the snapshot) is changed.

[Figure 1](#) shows the life cycle of a snapshot.

Figure 1. The snapshot life cycle



The following operations are applicable for the snapshot:

Create

Creates (takes) the 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.

- [Creating a snapshot](#)
 - [Locking and unlocking snapshots](#)
 - [Duplicating a snapshot](#)
 - [Creating a snapshot of a snapshot](#)
 - [Formatting a snapshot or a snapshot group](#)
- When duplicating a snapshot that has been changed using the unlock feature, the generated snapshot is actually a snapshot of a snapshot.
- This 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.

Throughout the volume life cycle, the data associated with the volume is continuously modified as part of the ongoing operation of the system. Whenever a request to modify a data portion on the master volume arrives, a copy of the original data is created and associated with the snapshot. Only then the volume is modified. This way, the data originally associated with the volume at the time the snapshot is taken is associated with the snapshot, effectively reflecting the way the data was before the modification.

Locking and unlocking snapshots

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 master 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.

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

Duplicating a snapshot

A user can create a new snapshot by duplicating an existing snapshot. The duplicate is identical to the source snapshot. The new snapshot is associated with the master 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 master volume, or the ability to restore from the snapshot.

Creating a snapshot of a snapshot

When duplicating a snapshot that has been 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 it is created, the new snapshot is viewed as another snapshot of the master volume.

Formatting a snapshot or a snapshot group

This 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 backup their volumes via snapshots, while maintaining the snapshot ID and the LUN ID. More than a single snapshot can be formatted per volume.

Required reading

Some of the concepts this topic refers to are introduced in this chapter as well as in a later chapter on this document. Consult the following reading list to get a grasp regarding these topics.

Snapshots

[Snapshot function and lifecycle](#)

Snapshot groups

[Consistency group snapshot lifecycle](#)

Attaching a host

[Host system attachment](#)

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 can't be unlocked.

No volume restore

The volume that the formatted snapshot belongs to can't be restored from it.

No restore from another snapshot

The formatted snapshot can't be restored from another snapshot.

No duplicating

The formatted snapshot can't be duplicated.

No re-format

The formatted snapshot can't be formatted again.

No volume copy

The formatted snapshot can't serve as a basis for volume copy.

No resize

The formatted snapshot can't be resized.

Use case

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

Restrictions in relation to other XIV operations

Snapshots of the following types can't be formatted:

Internal snapshot

Formatting an internal snapshot hampers the process it is part of, therefore is forbidden.

Part of a sync job

Formatting a snapshot that is part of a sync job renders the sync job meaningless, therefore is forbidden.

Part of a snapshot group

A snapshot that is part of a snapshot group can't 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 XCLI command that creates it. The snapshot's auto-generated name is derived from its volume's name and a serial number.

The following are examples of snapshot names:

MASTERVOL.snapshot_XXXXX
NewDB-server2.snapshot_00597

Parameter	Description	Example
MASTERVOL	The name of the volume.	NewDB-server2
XXXXX	A five-digit, zero filled snapshot number.	00597

Redirect-on-Write (ROW)

The IBM® XIV® 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 master volume:


1. The data originally associated with the master 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 master volume.

In contrast with the traditional copy-on-write 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.


In the following example of the Redirect-on-Write process, The volume is displayed with its data and the pointer to this data.

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




When a snapshot is taken, a new header is written first.

Figure 2. The 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.

Figure 3. The 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 keep updating by I/Os).

Figure 4. The 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 master volume. This master 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 master 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 master 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.



Storage pools

The storage space of the IBM® XIV® Storage System is portioned into *storage pools*, where each volume belongs to a specific storage pool.

Storage pools provide the following benefits:

Improved management of storage space

Specific volumes can be grouped together in a storage pool. This enables you to control the allocation of a specific storage space to a specific group of volumes. This storage pool can serve a specific group of applications, or the needs of a specific department.

Improved regulation of storage space

Snapshots can be automatically deleted when the storage capacity that is allocated for snapshots is fully consumed. This automatic deletion is performed independently on each storage pool. Therefore, when the size limit of the storage pool is reached, only the snapshots that reside in the affected storage pool are deleted. For more information, see [Additional snapshot attributes](#).

Facilitating thin provisioning

Thin provisioning is enabled by storage pools.

Storage pools as logical entities

A storage pool is a logical entity and is not associated with a specific disk or module. All storage pools are equally spread over all disks and all modules 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.

Note: For the size of the storage pool, please refer to the IBM XIV Storage System data sheet.

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

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 a storage pool level](#)

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

- [Thin provisioning](#)

The IBM XIV Storage System supports thin provisioning, which provides the ability to define logical volume sizes that are much larger than the physical capacity installed on the system. Physical capacity needs only to accommodate written data, while parts of the volume that have never been written to do not consume physical space.

Protecting snapshots at a storage pool level

Snapshots that participate 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 - 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 than this value.

Thin provisioning

The IBM® XIV® Storage System supports thin provisioning, which provides the ability to define logical volume sizes that are much larger than the physical capacity installed on the system. Physical capacity needs only to accommodate written data, while parts of the volume that have never been written to do not consume physical space.

Volume hard and soft sizes

Without thin provisioning, the size of each volume is both seen by the hosts and reserved on physical disks. Using thin provisioning, each volume is associated with the following two sizes:

Hard volume size

This reflects the total size of volume areas that were written by hosts. The hard volume size is not controlled directly by the user and depends only on application behavior. It starts from zero at volume creation or formatting and can reach the volume soft size when the entire volume has been written. Resizing of the volume does not affect the hard volume size.

Soft volume size

This is the logical volume size that is defined during volume creation or resizing operations. This is the size recognized by the hosts and is fully configurable by the user. The soft volume size is the traditional volume size used without thin provisioning.

System hard and soft size

Using thin provisioning, each IBM XIV Storage System is associated with a *hard system size* and *soft system size*. Without thin provisioning, these two are equal to the system's capacity. With thin provisioning, these concepts have the following meaning:

Hard system size

This is the physical disk capacity that was installed. Obviously, the system's hard capacity is an upper limit on the total hard capacity of all the volumes. The system's hard capacity can only change by installing new hardware components (disks and modules).

Soft system size

This is the total limit on the soft size of all volumes in the system. It can be set to be larger than the hard system size, up to 79TB. The soft system size is a purely logical limit, but should not be set to an arbitrary value. It must be possible to upgrade the system's hard size to be equal to the soft size, otherwise applications can run out of space. This requirement means that enough floor space should be reserved for future system hardware upgrades, and that the cooling and power infrastructure should be able to support these upgrades. Because of the complexity of these issues, the setting of the system's soft size can only be performed by IBM XIV support.

Pool hard and soft sizes

The concept of storage pool is also extended to thin provisioning. When thin provisioning is not used, storage pools are used to define capacity allocation for volumes. The storage pools control if and which snapshots are deleted when there is not enough space.

When thin provisioning is used, each storage pool has a soft pool size and a hard pool size, which are defined and used as follows:

Hard pool size

This is the physical storage capacity allocated to volumes and snapshots in the storage pool. The hard size of the storage pool limits the total of the hard volume sizes of all volumes in the storage pool and the total of all storage consumed by snapshots. Unlike volumes, the hard pool size is fully configured by the user.

Soft pool size

This is the limit on the total soft sizes of all the volumes in the storage pool. The soft pool size has no effect on snapshots.

Thin provisioning is managed for each storage pool independently. Each storage pool has its own soft size and hard size. Resources are allocated to volumes within this storage pool without any limitations imposed by other storage pools. This is a natural extension of the snapshot deletion mechanism, which is applied even without thin provisioning. Each storage pool has its own space, and snapshots within each storage pool are deleted when the storage pool runs out of space regardless of the situation in other storage pools.

The sum of all the soft sizes of all the storage pools is always the same as the system's soft size and the same applies to the hard size.

Storage pools provide a logical way to allocate storage resources per application or per groups of applications. With thin provisioning, this feature can be used to manage both the soft capacity and the hard capacity.

Depletion of hard capacity

Thin provisioning creates the potential risk of depleting the physical capacity. If a specific system has a hard size that is smaller than the soft size, the system will run out of capacity when applications write to all the storage space that is mapped to hosts. In such situations, the system behaves as follows:

Snapshot deletion

Snapshots are deleted to provide more physical space for volumes. The snapshot deletion is based on the deletion priority and creation time.

Volume locking

If all snapshots have been deleted and more physical capacity is still required, **all the volumes in the storage pool are locked and no write commands are allowed**. This halts any additional consumption of hard capacity.

Note: Space that is allocated to volumes that is unused (that is, the difference between the volume's soft and hard size) can be used by snapshots in the same storage pool.

The thin provisioning implementation in the IBM XIV Storage System manages space allocation per storage pool. Therefore, one storage pool cannot affect another storage pool. This scheme has the following advantages and disadvantages:

Storage pools are independent

Storage pools are independent in respect to the aspect of thin provisioning. **Thin provisioning volume locking on one storage pool does not create a problem in another storage pool.**

Space cannot be reused across storage pools

Even if a storage pool has free space, this free space is never reused for another storage pool. This creates a situation where volumes are locked due to the depletion of hard capacity in one storage pool, while there is available capacity in another storage pool.

Important: If a storage pool runs out of hard capacity, all of its volumes are locked to all write commands. Although write commands that overwrite existing data can be technically serviced, they are blocked to ensure consistency.

- **Instant space reclamation**

The IBM XIV Storage System instant space reclamation continuously recycles reusable IBM XIV storage space that is released by the host operating system, without any performance or management impact, and with measurable results.

Instant space reclamation

The IBM® XIV® Storage System instant space reclamation continuously recycles reusable IBM XIV storage space that is released by the host operating system, without any performance or management impact, and with measurable results.

Using instant space reclamation, storage and host administrators increase their systems capacity use and reduce the need for thin provisioning. Upon notification from the host, the IBM XIV frees any space that is no longer in use by writing zeroes into it. See more here: [Writing zeroes](#).

Communicating with the host in order to determine whether an allocated space is not in use evolves the following:

- Getting the allocation status from the host
- Matching this status with provisioning thresholds and reporting the findings
- Detecting space that is suitable for reclamation
- Freeing this space

The instant space reclamation feature skips volumes with temporary functionality and relations, such as:

- Off-line initialization of asynchronous mirroring
- Data migration of all kinds

The support of instant space reclamation for a mirrored pair of volumes is limited to synchronous mirroring and provided that both systems support the feature (that is, both systems are of version 11.2.0 and up).

Supported platforms

The following vendors have announced that by the end of 2012 their OS will support instant space reclamation:

- Microsoft Windows Server 8
- VMware
- RedHat
- Symantec SSF

Activating the instant space reclamation feature

Instant space reclamation can be globally enabled on manufacturing the IBM XIV Storage System, and disabled - and further enabled - by a technician.

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. Creating a synchronized snapshot set is especially important for applications that use multiple volumes concurrently. A typical example is a database application, where the database and the transaction logs reside on different storage volumes, but all of their snapshots must be taken at the same point in time.

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. Consistency group creation and options




All volumes in a consistency group must belong to the same storage pool.

- [Snapshot of a consistency group](#)
Taking a snapshot for the 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.
- [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 the 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



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 time stamp. This implies that the restored consistency group contains a consistent picture 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 master 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:

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.

QoS performance classes

The Quality of Service (QoS) feature allows the IBM® XIV® Storage System to deliver different service levels to hosts that are connected to the same XIV system.

The QoS feature favors performance of critical business applications that run concurrently with noncritical applications. Because the XIV 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 noncritical 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. The feature was extended in the XIV Storage Software Version 11.5 and can also be set by XIV domains and XIV storage pools. Each performance class is now implicitly one of two types: host type or pool/domain type.

The QoS feature possibilities and limitations can be summarized as follows:

- Up to 500 performance classes are configurable.
- QoS is applicable to host, domain, pool and restricted combinations of these entities. For instance, hosts cannot be specified for a performance class that already contains a domain or pool
- Limits can be defined as *Total*, meaning for XIV system as a whole, or *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 XIV interface nodes (equal to active interface modules). For example, on a 9-module partial rack XIV, where 4 interface modules are active, the total I/O or bandwidth rate would be divided by 4 (the number active interface modules). If a limit total of 3,000 I/Os is specified, it would result to a limitation of 750 I/Os per interface module.

In addition, in the case of the 9-module XIV, if the limited I/Os are coming through only two of the four interface modules (as a result of the SAN zoning), the effective limitation will be $2 \times 750 \text{ I/Os} = 1,500 \text{ I/Os}$ rather than the expected 3,000 I/O limitation.

Note: If more than one host, domain, or pool is added to a performance class, all hosts, domains, or pools in this performance class share the limitations defined on that performance class. For example, if two or more entities are added to a 10,000 IOPS performance class, the total number of all contained entities IOPS is limited to 10,000. Therefore, it is a good practice to create one performance class per domain and one performance class per 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 with hosts

The storage system connectivity is provided through the following interfaces:

- Fibre Channel for host-based I/O
- Gigabit Ethernet for host-based I/O using the iSCSI protocol
- Gigabit Ethernet for management (GUI or CLI) connectivity
- Remote access interfaces:
 - Call-home connection - connecting the IBM® XIV® Storage System to an IBM trouble-ticketing system.
 - Modem - for incoming calls only. The customer has to provide telephone line and number. The modem provides secondary means for providing remote access for IBM Support.

Figure 1. The IBM XIV Storage System interfaces



The following subsections provide information about different connectivity aspects.

- [IP and Ethernet connectivity](#)
- [Host system attachment](#)
The IBM XIV Storage System attaches to hosts of various operating systems.
- [Advanced host attachment](#)
The IBM XIV Storage System provides flexible host attachment options.
- [CHAP authentication of iSCSI hosts](#)
The MS-CHAP extension enables authentication of initiators (hosts) to the IBM XIV Storage System and vice versa in unsecured environments.
- [Clustering hosts into LUN maps](#)
To enhance the management of hosts, the IBM XIV Storage System allows clustering them together, where 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.
- [Supporting VMware extended operations](#)
The IBM XIV Storage System supports VMware extended operations that were introduced in VMware ESX Server 4 (VMware vStorage API).

IP and Ethernet connectivity

The following topics provide a basic explanation of the various Ethernet ports and IP interfaces that can be defined and various configurations that are possible within the IBM® XIV® Storage System.

The IBM XIV Storage System IP connectivity provides:

- iSCSI services over IP or Ethernet networks
- Management communication
- [Ethernet ports](#)
- [IPv6 certification](#)
The IBM XIV Storage System supports IPv6 and IPSec technology adoption as described in this topic.
- [Management connectivity](#)
- [Field technician ports](#)
- [Configuration guidelines summary](#)

Ethernet ports

The following three types of Ethernet ports are available:

iSCSI service ports

These ports are used for iSCSI over IP or Ethernet services. A fully equipped rack is configured with six Ethernet ports for iSCSI service. These ports should connect to the user's IP network and provide connectivity to the iSCSI hosts. The iSCSI ports can also accept management connections.

Management ports

These ports are dedicated for IBM® XIV® command-line interface (XCLI) and IBM XIV Storage Management GUI communications, as well as being used for outgoing SNMP and SMTP connections. A fully equipped rack contains three management ports.

Field technician ports

These ports are used for incoming management traffic only (usage is both XCLI and IBM XIV Storage Management GUI access). The ports are utilized only for the field technician's laptop computer and must not be connected to the user's IP network.

IPv6 certification

The IBM® XIV® Storage System supports IPv6 and IPSec technology adoption as described in this topic.

The IBM XIV Storage System supports IPv6 through stateless autoconfiguration and full IPSec (IKE2, transport, and tunnel mode) for Management and VPN ports.

Not supported

- There is no IPv6 support for technician notebook port.
- iSCSI ports are not supported.

Enabling and disabling IPv6

The IBM XIV Storage System supports IPv4 and IPv6 addresses out of the box. As the feature is enabled, stateless autoconfiguration is automatically enabled as well, and the system interfaces are getting ready to work with IPv6. Thus, looking for DNS addresses, the system also looks for AAAA entries.

Programs that are using connections on the Management and VPN ports must support IPv6 addresses. Each IP interface in the system may now 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.

The IPv6 addressing feature can be disabled if wanted.

Management connectivity

Management connectivity is used for the following functions:

- Executing XCLI commands through the IBM® XIV® command-line interface (XCLI)
- Controlling the IBM XIV Storage System through the IBM XIV Storage Management GUI
- Sending e-mail notification messages and SNMP traps about event alerts

To ensure management redundancy in case of module failure, the IBM XIV Storage System management function is accessible from three different IP addresses. 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 IBM XIV Storage Management GUI or XCLI for the set of 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.

XCLI and IBM XIV Storage Management GUI management

The IBM XIV Storage System management connectivity system allows users to manage the system from both the XCLI and IBM XIV Storage Management GUI. Accordingly, both can be configured to manage the system through iSCSI IP interfaces. Both XCLI and IBM XIV Storage Management GUI management is run over TCP port 7778. With all traffic encrypted through the Secure Sockets Layer (SSL) protocol.

System-initiated IP communication

The IBM XIV Storage System can also initiate IP communications to send event alerts as necessary. 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

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

Field technician ports

The IBM® XIV® Storage System supports two Ethernet ports. These ports are dedicated for the following reasons:

- Field technician use
- Initial system configuration
- Direct connection for service staff when they can not connect to customer network
- Directly manage the IBM XIV Storage System through a laptop computer

Laptop connectivity - configuring using DHCP

Two field technician ports are provided for redundancy purposes. This ensures that field technicians will always be able to connect a laptop to the IBM XIV Storage System. These two ports use a Dynamic Host Configuration Protocol (DHCP) server. The DHCP server will automatically assign IP addresses to the user's laptop and connects the laptop to the IBM XIV Storage System network. A laptop connected to any of the field technician ports is assigned an IP address and the IBM XIV Storage Management GUI or IBM XIV command-line interface (XCLI) will typically use the predefined configuration *direct-technician-port*.

Note: The two field technician laptop ports are used only to connect directly to the IBM XIV Storage System and should never be connected to the customer's network.

Laptop connectivity - configuring without DHCP

If the technician's laptop is not setup to receive automatic IP configuration information through DHCP, the laptop should be defined using these parameters:

IP address:
14.10.202.1
Netmask:
255.255.255.0
Gateway:
none
MTU:
1536

The field technician ports accept both XCLI and IBM XIV Storage Management GUI communications. SNMP and SMTP alerts are not sent through these ports.

Note: Each of the field technician ports is connected to a different module. Therefore, if a module fails, the port will become inoperative. When this happens, the laptop should be connected to the second port.

Configuration guidelines summary

When shipped, the IBM® XIV® Storage System does not have any IP management configurations. Accordingly, the following procedures should be performed when first setting up the system:

- Connecting a laptop to one of the field technician laptop ports on the patch panel
- Configuring at least one management IP interface
- Continuing the configuration process from either the technician port or from the configured IP interface

Note: It is important to define all three management IP interfaces and to test outgoing SNMP and SMTP connections from all three interfaces.

Host system attachment

The IBM® XIV® Storage System attaches to hosts of various operating systems.

The IBM XIV Storage System attaches to hosts through a set of Host Attachment Kits and complementary utilities.

Note: The term *host system attachment* was previously known as *host connectivity* or *mapping*. These terms are obsolete.

- **Balanced traffic without a single point of failure**
All host traffic (I/O) is served through up to six interface modules (modules 4-9). Although the IBM XIV 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.
- **Dynamic rate adaptation**
The IBM XIV Storage System provides a mechanism for handling insufficient bandwidth and external connections for the mirroring process.
- **Attaching volumes to hosts**
While the IBM XIV Storage System identifies volumes and snapshots by name, hosts identify volumes and snapshots according to their logical unit number (LUN).
- **Excluding LUN0**
Do not use LUN 0 as a normal LUN.

Balanced traffic without a single point of failure

All host traffic (I/O) is served through up to six interface modules (modules 4-9). Although the IBM® XIV® 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 watched and reviewed when host traffic patterns change. The IBM XIV 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 the host workload is adequately spread across the different connections and interface modules.

Dynamic rate adaptation

The IBM® XIV® Storage System provides a mechanism for handling insufficient bandwidth and external connections for the mirroring process.

The mirroring process replicates a local site on a remote site (See the [Synchronous remote mirroring](#) and [Asynchronous remote mirroring](#) chapters later on this document). To accomplish this, the process depends on the availability of bandwidth between the local and remote storage systems.

The mirroring process' sync rate attribute determines the bandwidth that is required for a successful mirroring. Manually configuring this attribute, the user takes into account the availability of bandwidth for the mirroring process, where the IBM XIV Storage System adjusts itself to the available bandwidth. Moreover, in some cases the bandwidth is sufficient, but external I/Os latency causes the mirroring process to fall behind incoming I/Os, thus to repeat replication jobs that were already carried out, and eventually to under-utilize the available bandwidth even if it was adequately allocated.

The IBM XIV Storage System prevents I/O timeouts through continuously measuring the I/O latency. Excess incoming I/Os are pre-queued until they can be submitted. The mirroring rate dynamically adapts to the number of pre-queued incoming I/Os, allowing for a smooth operation of the mirroring process.

Attaching volumes to hosts

While the IBM® XIV® Storage System identifies volumes and snapshots by name, 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.

Excluding LUN0

Do not use LUN 0 as a normal LUN.

LUN0 can be mapped to a volume just like other LUNs. However, when no volume is mapped to LUN0, the HAK is using it to discover the LUN array. Hence, we recommend not to use LUN0 as a normal LUN.

Advanced host attachment

The IBM® XIV® Storage System provides flexible host attachment options.

The following host attachment options are available:

- Definition of different volume mappings for different ports on the same host
- Support for hosts that have both Fibre Channel and iSCSI ports. Although it is not advisable to use these two protocols together to access the same volume, a dual configuration can be useful in the following cases:
 - As a way to smoothly migrate a host from Fibre Channel to iSCSI
 - As a way to access different volumes from the same host, but through different protocols

CHAP authentication of iSCSI hosts

The MS-CHAP extension enables authentication of initiators (hosts) to the IBM® XIV® Storage System and vice versa in unsecured environments.

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

Definitions

The following definitions apply to authentication procedures:

CHAP

Challenge Handshake Authentication Protocol

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.
Oneway (unidirectional CHAP)
CHAP authentication where initiators are authenticated by the target, but not vice versa.

Supported configurations

CHAP authentication type
Oneway (unidirectional) authentication mode, meaning that the Initiator (host) has to be authenticated by the IBM XIV Storage System.
MDS
CHAP authentication utilizes the MDS hashing algorithm.
Access scope
CHAP-authenticated Initiators are granted access to the IBM XIV Storage System via mapping that may restrict access to some volumes.

Authentication modes

The IBM XIV Storage System supports the following authentication modes:

None (default)
In this mode, an initiator is not authenticated by the IBM XIV Storage System.
CHAP (one way)
In this mode, an initiator is authenticated by the IBM XIV 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 IBM XIV storage system.

Changing the authentication mode from None to CHAP requires an authentication of the host. Changing the mode from CHAP to None doesn't require an authentication.

Complying with RFC 3720

The IBM XIV storage system CHAP authentication complies with the CHAP requirements [as stated in RFC 3720](#).


Secret length
The secret has to be between 96 bits and 128 bits; otherwise, the system fails the command, responding that the requirements are not fulfilled.
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 IBM® XIV® Storage System allows clustering them together, where 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.

Adding and removing hosts to a cluster are done as follows:

Figure 1. A volume, a LUN and clustered 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:

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

Removing a host from a cluster

The host is disbanded from the cluster, maintaining its connection to the LUN:


- The host's mapping remains identical to the mapping of the cluster.

- The mapping definitions do not revert to the host's original mapping (the mapping that was in effect before the host was added to the cluster).
- The host's mapping can be changed.

Notes:


- The IBM XIV Storage System defines the same mapping to all of the hosts of the same cluster. No hierarchy of clusters is maintained.
- Mapping a volume to a LUN that is already mapped to a volume.

Figure 2. You cannot map a volume to a LUN that is already mapped



- Mapping an already mapped volume to another LUN.

Figure 3. You cannot map a volume to a LUN, if the volume is already mapped.



• [Volume mapping exceptions](#)

The IBM XIV Storage System facilitates association of cluster mappings to a host that is added to a cluster. The system also facilitates easy specification of mapping exceptions for such host; such exceptions are warranted to accommodate cases where a host must have a mapping that is not defined for the cluster (e.g., Boot from SAN).

Volume mapping exceptions

The IBM® XIV® Storage System facilitates association of cluster mappings to a host that is added to a cluster. The system also facilitates easy specification of mapping exceptions for such host; such exceptions are warranted to accommodate cases where a host must have a mapping that is not defined for the cluster (e.g., Boot from SAN).

Mapping a volume to a host within a cluster

It is impossible to map a volume or a LUN that are already mapped.

For example, the host *host1* belongs to the cluster *cluster1* which has a mapping for the volume *vol1* to *lun1*:

1. Mapping *host1* to *vol1* and *lun1* **fails** as both volume and LUN are already mapped.
2. Mapping *host1* to *vol2* and *lun1* **fails** as the LUN is already mapped.
3. Mapping *host1* to *vol1* and *lun2* **fails** as the volume is already mapped.
4. Mapping *host1* to *vol2* and *lun2* **succeeds** with a warning that the mapping is host-specific.

Listing volumes that are mapped to a host/cluster

Mapped hosts that are part of a cluster are listed (that is, the list is at a host-level rather than cluster-level).

Listing mappings

For each host, the list indicates whether it belongs to a cluster.

Adding a host to a cluster

Previous mappings of the host are removed, reflecting the fact that the only relevant mapping to the host is the cluster's.
Removing a host from a cluster
The host regains its previous mappings.

Supporting VMware extended operations

The IBM® XIV® Storage System supports VMware extended operations that were introduced in VMware ESX Server 4 (VMware vStorage API).

The purpose of the VMware extended operations is to offload operations from the VMware server onto the storage system. The IBM XIV Storage System supports the following operations:

Full copy

The ability to copy data from one storage array to another without writing to the ESX server.

Block zeroing

Zeroing-out a block as a means for freeing it and make it available for provisioning.

Hardware-assisted locking

Allowing for locking volumes within an atomic command.

- **Writing zeroes**

The Write Zeroes command allows for zeroing large storage areas without sending the zeroes themselves.

- **Hardware-assisted locking**

The hardware-assisted locking feature utilizes VMware Compare and Write command for reading and writing the volume's metadata within a single operation.

- **Fast copy**

The Fast Copy functionality allows for VN cloning on the storage system without going through the ESX server.

Writing zeroes

The Write Zeroes command allows for zeroing large storage areas without sending the zeroes themselves.

Whenever a new VM is created, the ESXi server creates a huge file full of zeroes and sends it to the storage system. The Write Zeroes command is a way to tell a storage controller to zero large storage areas without sending the zeroes. To meet this goal, both VMware's generic driver and our own plug-in utilizes the WRITE SAME 16 command.

This method differs from the former method where the host used to write and send a huge file full of zeroes.

Note: The write zeroes operation is not a thin provisioning operation, as its purpose is not to allocate storage space.

Hardware-assisted locking

The hardware-assisted locking feature utilizes VMware Compare and Write command for reading and writing the volume's metadata within a single operation.

Upon the replacement of SCSI2 reservations mechanism with Compare and Write by VMware, the IBM® XIV® Storage System provides a faster way to change the metadata specific file, along with eliminating the necessity to lock all of the files during the metadata change.

The legacy VMware SCSI2 reservations mechanism is utilized whenever the VM server performs a management operation, that is to handle the volume's metadata. This method has several disadvantages, among them the mandatory overall lock of access to all volumes, which implies that all other servers are refrained from accessing their own files. In addition, the SCSI2 reservations mechanism entails performing at least four SCSI operations (reserve, read, write, release) in order to get the lock.

The introduction of the new SCSI command, called Compare and Write (SBC-3, revision 22), results with a faster mechanism that is displayed to the volume as an atomic action that does not require to lock any other volume.

Note: The IBM XIV Storage System supports single-block Compare and Write commands only. This restriction is carried out in accordance with VMware.

Backwards compatibility

The IBM XIV Storage System maintains its compatibility with older ESX versions as follows:

- Each volume is capable of connecting legacy hosts, as it still supports SCSI reservations.
- Whenever a volume is blocked by the legacy SCSI reservations mechanism, it is not available for an arriving COMPARE AND WRITE command.
- The Admin is expected to phase out legacy VM servers to fully benefit from the performance improvement rendered by the hardware-assisted locking feature.

Fast copy

The Fast Copy functionality allows for VN cloning on the storage system without going through the ESX server.

The Fast copy functionality speeds up the VM cloning operation by copying data inside the storage system, rather than issuing READ and WRITE requests from the host. This implementation provides a great improvement in performance, since it saves host to storage system intra-storage system communication. Instead, the functionality utilizes the huge bandwidth within the storage system.

IBM Real-time Compression with XIV

This section introduces the IBM® Real-time Compression™ (RtC) feature of IBM XIV storage systems.

IBM XIV storage system implementation of IBM RtC is a software-only feature that leverages the original XIV hardware design. IBM RtC, based on Random Access Compression Engine (RACE) technology, is field-proven since June 2012, starting with SVC and Storwize v7000 systems.

IBM Real-time Compression (RtC) was introduced with XIV storage system version 11.6 as an optional software feature for models 114 and 214. On the XIV Gen3 model 314, IBM RtC is available in the base license and enabled by default.

By doubling the XIV Gen3 RAM and CPU resources and dedicating the added resources to Real-time "Turbo" Compression, XIV model 314 further increases the high utilization and power efficiencies of XIV 4 TB and 6 TB disk drives, delivering outstanding data economics to your high-end storage.

Starting from version 11.6.1, you can uniformly and centrally manage all software licenses for storage deployments that are built with IBM Spectrum Accelerate software under one enterprise license agreement (ELA), including those for XIV Storage System Gen3. The application of Spectrum Accelerate software licenses to XIV storage system Gen3 hardware is supported.

Starting from version 11.6.2, XIV storage system also offers user-configurable soft capacity of up to 2 PB, and a reduced minimum compressible volume size from 103 GB (in version 11.6) to 51 GB.

IBM Real-time Compression with the IBM XIV storage system Gen3 effectively and efficiently answers a key requirement that typically challenges traditional approaches to data reduction. It reduces capacity while maintaining high performance of the storage system.

XIV implementation of IBM Real-time Compression highlights include:

- Substantial capacity savings across a versatile range of enterprise workloads.
- Use of the IBM Random Access Compression Engine (RACE) technology, which was purpose-built for real-life primary application workloads. IBM Real-time Compression takes advantage of data temporal locality to maximize data savings and system performance.
- Ease of use. An administrator can simply select the Compressed check box to create a new compressed volume. For an existing uncompressed volume, an accurate estimation of the potential compression savings is displayed in the XIV GUI. Assessing potential savings in the XIV GUI before compressing data requires considerably less time and effort than what it takes with external tools. Furthermore, non-disruptive conversion of uncompressed volumes to compressed volumes for existing volumes can provide an easy way to reclaim capacity and accelerate ROI.
Note: Compression is enabled by default on XIV Gen3 model 314 systems.
- Benefiting from the XIV architecture, compression compute resources are evenly distributed across the system, thereby increasing performance and efficiency.
- Due to the system's ability to preserve high performance consistency with compression, IBM Real-time Compression can be used with active primary data. Therefore, it supports workloads that are not candidates for compression in other solutions.
- [Planning for compression](#)
Before implementing IBM Real-time Compression on your system, assess the current types of data and volumes that are used on your system.

Turbo Compression in model 314

IBM XIV storage system Gen3 model 314 delivers Turbo Compression with larger effective capacity and guaranteed better performance:

- 2 x 6-core CPUs per module (versus 1 x 6-core CPU per module in model 214) - 1x 6-core CPU is dedicated to Real-time-Compression
- 96 GB RAM per module (versus 48 GB RAM per module in model 214) - 48 GB of RAM is dedicated to Real-time-Compression
- 1-2 PB of effective capacity without performance degradation
- Improved IOPS per compressed capacity

Benefits of IBM Real-time Compression

IBM Real-time Compression uses the reliable, field-proven, and patented IBM Random Access Compression Engine (RACE) technology to achieve a valuable combination of high performance and compression efficiencies. Data compression reduces the required storage capacity for a given amount of data, resulting in significantly lower Total Cost of Ownership (TCO).

Among the benefits of using IBM Real-time Compression are:

- Lower effective capacity requirements of a volume typically up to 1/5 of the uncompressed capacity.
- No additional hardware is required to use IBM Real-time Compression.
- Reduced cost for both software and hardware that is licensed by capacity because less physical storage is required for compressed data.
- If you already have 214, save Capital Expenditure (CAPEX) by only purchasing an IBM XIV RtC software license, leverage your current XIV Gen3 investment, and apply compression to existing data and new data.
- Compression is transparent to the applications and can be enabled or disabled on any volume, at any time, non-disruptively.
- Compressed volumes can be mirrored like other XIV volumes. The required bandwidth for a compressed volume is significantly reduced since the replicated data is compressed.

For the same reason, mirroring and Hyper-Scale Mobility are faster and require less bandwidth because less data is transferred.

Remote volume copies are always compressed if the source is compressed.

Planning for compression

Before implementing IBM Real-time Compression on your system, assess the current types of data and volumes that are used on your system.

Note: If you are using IBM XIV storage system software version 11.6.x or later, and IBM XIV Management Tools version 4.6 or later, you can see storage saving estimates (even if compression is not licensed or is just disabled). The decision to compress volumes can be based on the expected storage savings of the compressed data and the expected effect on performance.

- [Understanding compression rates, ratios and savings](#)
Consider a use case where the original capacity required to hold data was 100 TB, but 20 TB after compression (100 TB = 20 TB + 80 TB saved).
- [Prerequisites and limitations](#)
The following prerequisites and limitations are for XIV storage system models 114 or 214, with XIV Storage software version 11.6 or later, and models 214 or 314 with XIV storage software version 11.6.1 or later.
- [Estimating compression savings](#)
Compressible data can be identified and expected compression ratios can be estimated even before using compression.

Understanding compression rates, ratios and savings

Consider a use case where the original capacity required to hold data was 100 TB, but 20 TB after compression (100 TB = 20 TB + 80 TB saved).

The following values help to clarify these terms:

Compression rate = 80%

Compression savings rate = 80%

Compression savings = 80 TB

Compression ratio = original size (100 TB) *divided by* the size on disk after compression (20 TB) = 5:1

When you consider savings, it is easiest to use the compression rate.

The compression ratio helps in understanding how much effective data you can store on your system. So, when you have a 5:1 compression ratio, you will be able to store 500 TB of data on 100 TB of physical capacity.

Prerequisites and limitations

The following prerequisites and limitations are for XIV storage system models 114 or 214, with XIV Storage software version 11.6 or later, and models 214 or 314 with XIV storage software version 11.6.1 or later.

Tip: Consider using version 11.6.1 with model 214 to benefit from improvements like compressed volume size.

- Compressed volumes must be created in thin-provisioned pools.
- To convert a mirrored volume from uncompressed to compressed (or vice versa), the mirroring relationship for that volume must first be removed and then recreated after the conversion.
- Snapshots cannot be converted from compressed to uncompressed and vice versa. Snapshots that already existed before a volume was converted from non-compressed to compressed are not converted and are not available with the converted volume.
- Space requirements:
 - Prior to enabling compression, the system must have a minimum of 17 GB of free hard space available. Enabling compression reserves 17 GB from the available system hard capacity. It is reserved for internal system use only.
 - Before the compression process, there must be enough space for both compressed and uncompressed versions of the volume.
 - For models 114 and 214 with version 11.6.0: Volume size must be at least 103 GB before compression.
 - For models 214 and 314 with version 11.6.1: Volume size must be at least 51 GB before compression.

The following is a partial list of limitations.

- Up to 1024 volumes and snapshots can be compressed.
- The following limits apply to compression capacity:
 - System must have a minimum of 17 GB of free hard space to enable IBM Real-time Compression.
 - Thin pools require a minimum of 17 GB of free hard space available to convert or transform volumes from uncompressed to compressed.
 - Thin pools require a minimum free soft space that is at least as large as the volume size that is being converted from uncompressed to compressed.
 - When you decompress a compressed volume, you must have both free hard space at least the size of the uncompressed volume and free soft space. It is a good practice to have free soft space at least the size of the uncompressed volume.
 - The Storage Admin can modify the system soft capacity.
Tip: Over-provisioning with Real-time Compression is safe, since compression ratios are predictable and stable.
- To compress an uncompressed volume in a thick pool, it must be moved to a thin-provisioned pool with compression enabled.
- Only one conversion process can be active at any time.
- Adding a module, rebuilding a disk, or upgrading the system suspends and then resumes the conversion process.

For the most current information about the limitations, refer to the [Limitations](#) section of the IBM XIV storage system Gen3 Release Notes, versions 11.6.0, 11.6.1 and 11.6.2.

Estimating compression savings

Compressible data can be identified and expected compression ratios can be estimated even before using compression.

On an XIV system supporting compression, the compression ratio for all uncompressed volumes in the system is continuously estimated, even before enabling compression. The decision to use compression can be based on the expected storage savings of the compressed data and the expected effect on performance (throughput and latency) of the compression processing overhead.

Information on compression usage can also be monitored using the XIV GUI to determine the potential savings to your storage capacity when uncompressed volumes are compressed. You can view the total percentage and total size of capacity savings when compression is used on the system. Compression savings across individual domains, pools and volumes can also be monitored. These compression values can be used to determine which volumes have achieved the highest compression savings. See the *IBM® XIV® Management Tools Operations Guide* for more information on monitoring and using compression.

Note: Keep in mind that the expected storage savings can vary from 5% higher or lower than the actual compression ratio. A negative estimated compression ratio can be due to metadata that consumes storage space on a volume, even when a 0% estimated compression ratio is received from data that cannot be compressed.

- [Effective capacity](#)
Effective capacity is the amount of storage that is virtually allocated to applications.
- [General compression saving guidelines](#)
The best candidates for data compression are data types that are not already compressed. Compressed data types are used by many workloads and applications, such as databases, character/ASCII-based data, email systems, server virtualization, CAD/CAM, software development systems, and vector data.
- [Estimating compression savings using XIV GUI](#)
The XIV software provides a built-in comprestimator function from the XIV Management Tools GUI on an XIV system supporting Real-time Compression.
- [Estimating compression savings using IBM Comprestimator utility](#)
Comprestimator is a stand-alone tool that can be used to estimate compression savings for data that is either not on XIV storage, or on an XIV Gen2 or Gen3 storage system with system software earlier than 11.6.x.

Effective capacity

Effective capacity is the amount of storage that is virtually allocated to applications.

Using thin-provisioned storage architectures, the effective capacity is larger than the array usable capacity. This is made possible by over-committing capacity, or by compressing the served data. Compression is the preferred method to apply thin-provisioning to usable capacity, since the compression ratio is highly predictable.

Hard capacity denotes usable, non-compressed capacity, whereas *soft capacity* denotes the nominal capacity that is assigned to volumes, and reported to any hosts mapped to those volumes. Thin provisioning denotes committing more soft capacity than hard capacity. Soft capacity is assigned at a pool level. In the case of compression, thin-provisioning is obvious: a compressed volume will forever use less hard capacity than soft capacity.

In XIV terms, the effective capacity is allocated out of the system soft capacity, and is the sum of the sizes of all the allocated volumes.

In an XIV System Gen3 Turbo Compression model 314, the maximal, non-compressed hard capacity supported in a single XIV frame is 485 TB (15 modules, 6 TB drives). However, an XIV frame can effectively accommodate up to 2 PB of real written data - when the data is compressed. With very high compression rates, filling a system up to 2 PB of soft capacity may not require a lot of usable capacity.

The maximum *soft capacity* that can be allocated to volumes in XIV is 2 PB. Considering the compression ratio for typical data profiles on XIV systems, the effective soft capacity leveraged within a 15-module frame will range from 1 PB to 2 PB. To maximize the utilization of XIV hard and soft capacity with compressed data, and avoid over-sizing the system, it is important to assess the expected compression ratio for the stored data. The Comprestimator command-line host-based utility can be used for that purpose. For more information on Comprestimator, see [Estimating compression savings using IBM Comprestimator utility](#).

The maximum effective capacity is reached when all the soft capacity has been allocated to volumes (that is, 2 PB).

For more information on how to fine-tune soft and hard capacities, refer to "Additional space utilization guidance" in the [Real-time Compression with IBM XIV Storage System Model 314 \(REDP-5306\)](#) Redbook.

General compression saving guidelines

The best candidates for data compression are data types that are not already compressed. Compressed data types are used by many workloads and applications, such as databases, character/ASCII-based data, email systems, server virtualization, CAD/CAM, software development systems, and vector data.

The following examples represent workloads and data that are already compressed and are, therefore, not good candidates for compression.

- Compressed audio, video, and image file formats -
File types such as JPEG, PNG, MP3, medical imaging (DICOM), and MPEG2
- Compressed user productivity file formats -
Microsoft Office 2007 newer formats (.pptx, .docx, .xlsx, and so on), PDF files, Microsoft Windows executable files (.exe), and so on
- Compressed file formats
File types such as .zip, .gzip, .rar, .cab, and .tgz


Note: The required capacity reserve is equal to the size of the volume (not the used capacity, but the volume size).

Estimating compression savings using XIV GUI

The decision to use compression can be based on the expected storage savings of the compressed data and the expected effect on performance (throughput and latency) of the compression processing overhead.

Figure 1 displays the compression savings (in both percentage and GB values) of compressed volumes and uncompressed volumes with *estimates* of potential savings, should the uncompressed volumes be compressed. These potential compression savings are constantly being updated.

Figure 1. Compression savings in the Volumes by Pools view



Compression Saving and *Compression Saving (%)* appear on the following views:

- Storage Pools
- Volumes by Pools
- Volumes and Snapshots
- Consistency Groups
- Domains
- Systems list

Estimating compression savings using IBM Comprestimator utility

Comprestimator is a stand-alone tool that can be used to estimate compression savings for data that is either not on XIV storage, or on an XIV Gen2 or Gen3 storage system with system software earlier than 11.6.x.

Comprestimator is a command-line host-based utility that can be used to estimate the expected compression rate for block-devices. The utility uses advanced mathematical and statistical algorithms to perform sampling and analysis efficiently. The utility also displays its accuracy level by showing the compression accuracy range of the results that are achieved based on the formulas it uses, deviating plus or minus 5 percent based on the formulas that are used by the RACE implementation.

The utility runs on a host that has access to the devices to be analyzed. It runs only read operations, so it has no effect on the data that is stored on the device. The following links provide useful information about installing Comprestimator on a host and using it to analyze devices on that host: [Comprestimator Utility](#) and [Comprestimator Host](#)

Synchronous remote mirroring

IBM® XIV® Storage System features synchronous and asynchronous remote mirroring for disaster recovery. Remote mirroring can be used to replicate the data between two geographically remote sites. The replication ensures uninterrupted business operation if there is a total site failure.

Remote mirroring provides data protection for the following types of site disasters:

Site failure

When a disaster happens to a site that is remotely connected to another site, the second site takes over and maintains full service to the hosts connected to the first site. The mirror is resumed after the failing site recovers.

Split brain

After a communication loss between the two sites, each site maintains full service to the hosts. After the connection is resumed, the sites complement each other's data to regain mirroring.

Synchronous and asynchronous remote mirroring

The two distinct methods of remote mirroring - synchronous and asynchronous - are described on this chapter and on the following chapter. Throughout this chapter, the term *remote mirroring* refers to *synchronous* remote mirroring, unless clearly stated otherwise.

- [Remote mirroring basic concepts](#)
Synchronous remote mirroring provides continuous availability of critical information in the case of a disaster scenario.
- [Remote mirroring operation](#)
Remote mirroring operations involve *configuration, initialization, ongoing operation, handling of communication failures, and role switching* activities.
- [Configuration options](#)
The remote mirroring configuration process involves configuring volumes and volume pair options.
- [Synchronous mirroring statuses](#)
The status of the synchronous remote mirroring volume represents the state of the storage volume in regard to its remote mirroring operation.
- [I/O operations](#)
I/O operations are performed on the primary and secondary volumes across various configuration options.
- [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.
- [Role switchover](#)
- [Remote mirroring](#)

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 master storage system.
A local site that contains both the data and the active servers.

Secondary site

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

Master

The volume or storage system which is mirrored. The master volume or storage system is usually at the master site.

Slave

The volume or storage system to which the master is mirrored. The slave volume or storage system is usually at the Secondary site.

One of the main goals of remote mirroring is to ensure that the secondary site contains the same (consistent) data as the master site. With remote mirroring, services are provided seamlessly by the hosts and storage systems at the secondary site.

The process of ensuring that both storage systems contain identical data at all times is called *remote mirroring*. Synchronous remote mirroring is performed during each write operation. The write operation issued by a host is sent to both the master and the slave storage systems.

To ensure that the data is also written to the secondary system, acknowledgment of the write operation is only issued after the data has been written to both storage systems. This ensures the consistency of the secondary storage system to the master storage system except for the contents of any last, unacknowledged write operations. This form of mirroring is called synchronous mirroring.

In a remote mirroring system, reading is performed from the master storage system, while writing is performed on both the master and the slave storage systems, as previously described.

The IBM® XIV® Storage System supports configurations where server pairs perform alternate master or secondary roles with respect to their hosts. As a result, a server at one site might serve as the master storage system for a specific application, while simultaneously serving as the secondary storage system for another

application.

Remote mirroring operation

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

The following list defines the remote mirroring operation activities:

Configuration

Local and remote replication peers are defined by an administrator who specifies the primary and secondary volumes. For each coupling, several configuration options can be defined.

Initialization

Remote mirroring operations begin with a master volume that contains data and a formatted slave volume. The first step is to copy the data from the master volume to the slave volume. This process is called *initialization*. Initialization is performed once in the lifetime of a remote mirroring coupling. After it is performed, both volumes are considered synchronized.

Ongoing Operation

After the initialization process is complete, remote mirroring is activated. During this activity, all data is written to the master volume and then to the slave volume. The write operation is complete after an acknowledgment from the slave volume. At any point, the master and slave volumes are identical except for any unacknowledged (pending) writes.

Handling of Communication Failures

From time to time the communication between the sites might break down, and it is usually preferable for the primary site to continue its function and to update the secondary site when communication resumes. This process is called *synchronization*.

Role Switching

When needed, a replication peer can change its role from master to slave or vice versa, either as a result of a disaster at the primary site, maintenance operations, or because of a drill that tests the disaster recovery procedures.

Configuration options

The remote mirroring configuration process involves configuring volumes and volume pair options.

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

Table 1. Configuration options for a volume

Name	Values	Definition
Role	None, Master, Slave	Role of a volume. (Primary and Secondary are designations.)
Peer	Remote target identification and the name of the volume on the remote target.	Identifies the peer volume.

Table 2. Configuration options for a coupling

Name	Values	Definition
Activation	Active, Stand-by.	Activates or deactivates remote mirroring.

- [Volume configuration](#)
The role of each volume and its peer volumes on the IBM XIV Storage System must be defined for it to function within the remote mirroring process.
- [Communication errors](#)
- [Coupling activation](#)
Remote mirroring can be manually activated and deactivated per coupling. When it is activated, the coupling is in *Active* mode. When it is deactivated, the coupling is in *Standby* mode.

Volume configuration

The role of each volume and its peer volumes on the IBM® XIV® Storage System must be defined for it to function within the remote mirroring process.

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

- Volume role
- Peer

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 configured as part of any remote mirroring configuration.

Master volume

The volume is part of a mirroring coupling and serves as the master volume.

All write operations are made to this master volume. It ensures that write operations are made to the slave volume before acknowledging their success.

Slave volume

This volume is part of a mirroring coupling and serves as a backup to the master volume.

Data is read from the slave volume, but cannot be written to it.

A *peer* is a volume that is part of a coupling. A volume with a role other than none has to have a peer designation, and a corresponding master or slave volume assigned to it.

Configuration errors

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.

- **Mixed configuration**
The volumes on a single storage system can be defined in a mixture of configurations.

Mixed configuration

The volumes on a single storage system can be defined in a mixture of configurations.

For example, a storage system can contain volumes whose role is defined as master, as well as volumes whose roles are defined as slave. In addition, some volumes might not be involved in a remote mirroring coupling at all.

The roles assigned to volumes are *transient*. This means a volume that is currently a master volume can be defined as a slave volume and vice versa. The term *local* refers to the master volume, and *remote* refers to the slave volume for processes that switch the master and slave assignments.

Communication errors

When the communication link to the secondary volume fails or the secondary volume itself is not usable, processing to the volume continues as usual. The following occurs:

- The system is set to an unsynchronized state.
- All changes to the master volume are recorded and then applied to the slave volume after communication is restored.

Coupling activation

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

These modes have the following functions:

Active

Remote mirroring is functioning and the data is being written to both the master and the slave volumes.

Standby

Remote mirroring is deactivated. The data is not being written to the slave volume, but it is being recorded in the master volumes which will later synchronize the slave 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 master volumes do not generate alerts that the mirroring has failed.

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.
- Transitions between the two states can only be performed from the UI and on the volume.

Synchronous mirroring statuses

The status of the synchronous remote mirroring volume represents the state of the storage volume in regard to its remote mirroring operation.

The state of the volume is a function of the status of the communication link and the status of the coupling between the master volume and the slave volume. [Link status](#) describes the various statuses of a synchronous remote mirroring volume during remote mirroring operations.

Table 1. Synchronous mirroring statuses

Entity	Name	Values	Definition
Link	Status	• Up • Down	Specifies if the communications link is up or down.
Coupling	Operational status	• Operational • Non-operational	Specifies if remote mirroring is working.

Entity	Name	Values	Definition
	Synchronization status	• Initialization • Synchronized • Unsynchronized • Consistent • Inconsistent	Specifies if the master and slave volumes are consistent.
	Last-secondary-timestamp	Point-in-time date	Time stamp for when the secondary volume was last synchronized.
	Synchronization process progress	Synchronization status	Amount of data remaining to be synchronized between the master and slave volumes due to non-operational coupling.
	Secondary locked	Boolean	True, if secondary was locked for writing due to lack of space; otherwise false. This can happen during the synchronization process when there is not enough space for the last-consistent snapshot.
	Configuration error	Boolean	True, if the configuration of the master and secondary slave is inconsistent.

- [Link status](#)
The status of the communication link can be either *up* or *down*. The link status of the master volume is, of course, also the link status of the slave volume.
- [Operational status](#)
The coupling between the master and slave volumes is either *operational* or *non-operational*. 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 operational status is non-operational.
- [Synchronization status](#)
The synchronization status reflects the consistency of the data between the master and slave volumes. Because the purpose of the remote mirroring feature is to ensure that the slave volume is an identical copy of the master volume, this status indicates whether this objective is currently attained.

Link status

The status of the communication link can be either *up* or *down*. The link status of the master volume is, of course, also the link status of the slave volume.

Operational status

The coupling between the master and slave volumes is either *operational* or *non-operational*. 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 operational status is non-operational.

Synchronization status

The synchronization status reflects the consistency of the data between the master and slave volumes. Because the purpose of the remote mirroring feature is to ensure that the slave volume is an identical copy of the master volume, this status indicates whether this objective is currently attained.

The possible synchronization statuses for the master volume are:

Initialization

The first step in remote mirroring is to create a copy of the data from the master volume to the slave volume, at the time when the mirroring was set to place. During this step, the coupling status remains initialization.

Synchronized (master 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 written to its peer volume. This means that their write operations have not yet been acknowledged. 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 written 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.

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.

- **[Synchronization progress status](#)**
During the synchronization process while the secondary volumes are being updated with previously written data, the volumes have a dynamic synchronization process status.
- **[Last secondary timestamp](#)**
A timestamp is taken when the coupling between the primary and secondary volumes becomes non-operational.

Synchronization progress status

During the synchronization process while the secondary volumes are being updated with previously written data, the volumes have a dynamic synchronization process status.

This status is comprised of 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

Estimate of the time, which is required to complete the Synchronization process and achieve synchronization, based on past rate.

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.

I/O operations

I/O operations are performed on the primary and secondary volumes across various configuration options.

I/O on the primary

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, an error is returned.

The error reflects the type of problem that was encountered. For example, remote mirroring has been deactivated, there is a locked secondary error, or there is a link error.

I/O on the secondary

A secondary volume can have LUN maps and hosts associated with it, but it is only accessible as a read-only volume. These maps are used by the backup hosts when a switchover is performed. When the secondary volume becomes the primary volume, 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 the new primary volume.

Read

Data is read from the secondary volume like from any other volume.

Write

An attempt to write on the secondary volume results in a volume read-only SCSI error.

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.


- **State diagram**
Couplings can be in the *Initialization*, *Synchronized*, *Timestamp*, or *Unsynchronized* state.
- **Coupling recovery**
Remote mirroring recovers from non-operational coupling.
- **Uncommitted data**
When the coupling is in an Unsynchronized state, for best-effort coupling, 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.
- **Constraints and limitations**
Coupling has constraints and limitations.
- **Last-consistent snapshots**
Before the synchronization process is initiated, a snapshot of the secondary volume is created. This snapshot is created to ensure the usability of the secondary volume in case of a primary site disaster during the synchronization process.
- **Secondary locked error status**
While the synchronization process is in progress, there is a period in which the secondary volume is not consistent with the primary volume, and a last-consistent snapshot is maintained. While in this state, I/O operations to the secondary volume can fail because there is not enough space. There is not enough space because every I/O operation potentially requires a copy-on-write of a partition.

State diagram

Couplings can be in the *Initialization*, *Synchronized*, *Timestamp*, or *Unsynchronized* state.

The following diagram shows the various coupling states that the IBM® XIV® Storage System assumes during its lifetime, along 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 both the primary and secondary volumes are consistent.

Timestamp

Remote mirroring has become non-operational so a time stamp is recorded. During this status, the following actions take place:

1. Coupling deactivation, or the link is down
2. Coupling reactivation, or the link is restored.

Unsynchronized

Remote mirroring is non-operational because of a communications failure or because remote mirroring was deactivated. Therefore, the primary and secondary volumes are not synchronized. When remote mirroring resumes, steps are taken to return to the Synchronized state.

Coupling recovery

Remote mirroring recovers from non-operational coupling.

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.
- The primary volume deletes the special snapshot after all couplings that mirror volumes between the same pair of systems are synchronized.

Uncommitted data

When the coupling is in an Unsynchronized state, for best-effort coupling, 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*.

Note: There is only metadata that marks the parts of the primary volume that must be written to the secondary volume when the coupling becomes operational.

Constraints and limitations

Coupling has constraints and limitations.

The following constraints and limitations exist:

- The Size, Part, or Time-to-synchronize are relevant only if the Synchronization status is Unsynchronized.
- The last-secondary-time stamp is only relevant if the coupling is Unsynchronized.

Last-consistent snapshots

Before the synchronization process is initiated, a snapshot of the secondary volume is created. This snapshot is created to ensure the usability of the secondary volume in case of a primary site disaster during the synchronization process.

If the primary volume is destroyed before the synchronization is completed, the secondary volume might be inconsistent because it may have been only partially updated with the changes that were made to the primary volume. The reason for this possible inconsistency is the fact that the updates were not necessarily performed in the same order in which they were written by the hosts.

To handle this situation, the primary volume always creates a snapshot of the last-consistent secondary volume after re-connecting to the secondary machine, and before starting the synchronization process.

The last consistent snapshot

The Last Consistent snapshot (LCS) is created by the system on the Slave peer in synchronous mirroring just before mirroring resynchronization needs to take place. Mirroring resynchronization takes place after link disruption, or a manual mirroring deactivation. In both cases the Master will continue to accept host writes, yet will not replicate them onto the Slave as long as the link is down, or the mirroring is deactivated.

Once the mirroring is restored and activated, the system takes a snapshot of the Slave (LCS), which represents the data that is known to be mirrored, and only then the non yet mirrored data written to the Master is replicated onto the Slave through a resynchronization process.

The LCS is deleted automatically by the system once the resynchronization is complete for all mirrors on the same target, but if the Slave peer role is changed during resynchronization - this snapshot will not be deleted.

The external last consistent snapshot

Prior to the introduction of the external last consistent snapshot, whenever the peer's role was changed back to Slave and sometime whenever a new resynchronization process had started, the system would detect an LCS on the peer and would not create a new one. If, during such an event, the peer was not part of a mirrored consistency group (mirrored CG) it would mean that not all volumes have the same LCS timestamp. If the peer was part of a mirrored consistency group, we would have a consistent LCS but not as current as possibly expected. This situation is avoided thanks to the introduction of the external last consistent snapshot.

Whenever the role of a Slave with an LCS is changed to Master while mirroring resynchronization is in progress (in the system/target not specific to this volume), the LCS is renamed external last consistent (ELCS). The ELCS retains the LCS deletion priority of 0. If the peer's role is later changed back to Slave and sometime afterwards a new resynchronization process starts, a new LCS will be created.

Subsequently changing the Slave role again will rename the existing external last consistent snapshot to external last consistent x (x being the first available number starting from 1) and will rename the LCS to external last consistent. The deletion priority of external last consistent will be 0, but the deletion priority of the new external last consistent x will be the system default (1), and can thus be deleted automatically by the system upon pool space depletion.

It is crucial to validate whether the LCS or an ELCS (or even ELC x) should serve as a restore point for the Slave peer if resynchronization cannot be completed. While snapshots with deletion priority 0 are not automatically deleted by the system to free space, the external last consistent and external last consistent x snapshots can be manually deleted by the administrator if so required. As the deletion of such snapshots might leave an inconsistent peer without a consistent snapshot to be restored from (in case the resynchronization cannot complete due to Master unavailability) - it should generally be avoided even when pool space is depleted, unless the peer is guaranteed to be consistent.

Manually deleting the last consistent snapshot

- Only the XIV support team can delete the last consistent snapshot.
- The XIV support team can also configure a mirroring so that it does not create the last consistent snapshot. This is required when the system that contains the secondary volume is fully utilized and an additional snapshot cannot be created.

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

This table provides an example of a failure situation and describes the time specified by the timestamp.

Table 1. Example of the last consistent snapshot time stamp process

Time	Status of coupling	Action	Last consistent timestamp
Up to 12:00	Operational and synchronized		Current
12:00 - 1:00	Failure caused the coupling to become non-operational. The coupling is Unsynchronized.	Writing continues to the primary volume. Changes are marked so that they can be committed later.	12:00
13:00	Connectivity resumes and remote mirroring is operational. Synchronization begins. The coupling is still Unsynchronized.	All changes since the connection was broken are committed to the secondary volume, as well as current write operations.	12:00
13:15	Synchronized		Current

Secondary locked error status

While the synchronization process is in progress, there is a period in which the secondary volume is not consistent with the primary volume, and a last-consistent snapshot is maintained. While in this state, I/O operations to the secondary volume can fail because there is not enough space. There is not enough space because every I/O operation potentially requires a copy-on-write of a partition.

Whenever I/O operations fail because there is not enough space, all couplings in the system are set to the secondary-locked status and the coupling becomes non-operational. The administrator is notified of a critical event, and can free space on the system containing the secondary volume.

If this situation occurs, contact an IBM XIV field engineer. After there is enough space, I/O operations resume and remote mirroring can be reactivated.

Role switchover

- [Role switchover when remote mirroring is operational](#)
- [Role switchover when remote mirroring is nonoperational](#)
- [Resumption of remote mirroring after role change](#)

Role switchover when remote mirroring is operational

Role switching between primary and secondary volumes can be performed from the IBM® XIV® Storage Management GUI or the XCLI after the remote mirroring function is operational. After role switchover occurs, the primary volume becomes the secondary volume and vice versa.

There are two typical reasons for performing a switchover when communications between the volumes exist. These are:

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.

This switchover is performed as an automatic operation acting on the primary volume. This switchover cannot be performed if the primary and secondary volumes are not synchronized.

Role switchover when remote mirroring is nonoperational

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.

Switchover procedures differ depending on whether the primary and secondary volumes are connected or not. As a general rule, the following is true:

- When the coupling is deactivated, it is possible 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 nonoperational 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.

- [Switch secondary to primary](#)
 - [Secondary consistency](#)
 - Switching the secondary volume to primary, when the last-consistent snapshot is no longer available
 - [Switch primary to secondary](#)
- When coupling is inactive, the primary machine can switch roles. After such a switch, the primary volume becomes the secondary one.

Switch secondary to primary

The role of the secondary volume can be switched to primary using the IBM® XIV® Storage Management GUI or the XCLI. After this switchover, the following is true:

- The secondary volume is now the primary.
- 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 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.

Secondary consistency

Switching the secondary volume to primary, when the last-consistent snapshot is no longer available

If the user is switching the secondary to a primary volume, and a snapshot of the last_consistent state exists, then the link was broken during the process of synchronizing. In this case, the user has a choice between using the most-updated version, which might be inconsistent, or reverting to the last_consistent snapshot. [Table 1](#) outlines this process.

Table 1. Disaster scenario leading to a secondary consistency decision

Time	Status of remote mirroring	Action
Up to 12:00	Operational	Volume A is the primary volume and volume B is the secondary volume.
12:00	Non-operational because of communications failure	Writing continues to volume A and volume A maintains the log of changes to be committed to volume B.
13:00	Connectivity resumes and remote mirroring is operational	A last_consistent snapshot is generated on volume B. After that, volume A starts to update volume B with the write operations that occurred since communication was broken.
13:05	Primary site is destroyed and all information is lost	
13:10		Volume B is becoming the primary. The operators can choose between using the most-updated version of volume B, which contains only parts of the write operations committed to volume A between 12:00 and 13:00, or use the last-consistent snapshot, which reflects the state of volume B at 12:00.

If a last-consistent snapshot exists and the role is changed from secondary to primary, the system automatically generates a snapshot of the volume. This snapshot is named *most_updated snapshot*. It is generated to enable a safe reversion to the latest version of the volume, when recovering from user errors. This snapshot can only be deleted by the IBM® XIV® Storage System support team.

If the coupling is still in the initialization state, switching cannot be performed. In the extreme case where the data is needed even though the initial copy was not completed, a volume copy can be used on the primary volume.

Switch primary to secondary

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

Because the primary volume is inactive, it is also in the unsynchronized state, and it might have an uncommitted data list. All these changes will be lost. When the volume becomes secondary, this data must be reverted 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 changes that were lost.

The uncommitted data list has now switched its semantics, and instead of being a list of updates that the local volume (old primary, new secondary) needs to update on the remote volume (old secondary, new primary), the list now represents the updates that need to be taken from the remote to the local volume.

Upon reestablishing the connection, the local 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).

Resumption of remote mirroring after role change

When the communication link is resumed after a switchover of roles in which both sides were switched, the coupling now contains one secondary and one primary volume.

Note: After a role switchover, the coupling is in standby. The coupling can be activated before or after the link resumes.

[Table 1](#) describes the system when the coupling becomes operational, meaning after the communications link has been resumed and the coupling has been reactivated. When communications is resumed, the new primary volume (old secondary) might be in the unsynchronized state, and have an uncommitted data list to synchronize.

The new secondary volume (old primary), might have an uncommitted data list to synchronize from the new primary volume. These are write operations that were written after the link was broken and before the role of the volume was switched from primary to secondary. These changes must be reverted to the content of the new primary volume. Both lists must be used for synchronization by the new primary volume.

Table 1. Resolution of uncommitted data for synchronization of the new primary volume

Time	Status of remote mirroring	Action
Up to 12:00	Operational and synchronized	Volume A is the primary volume and volume B is the secondary volume.
12:00 to 12:30	Communication failure, coupling becomes non-operational	Volume A still accepts write operations from the hosts and maintains an uncommitted data list marking these write operations. For example, volume A accepted a write operation to blocks 1000 through 2000, and marks blocks 1000 through 2000 as ones that need to be copied to volume B after reconnection.
12:30	Roles changed on both sides	Volume A is now secondary and volume B is primary. Volume A should now revert the changes done between 12:00 and 12:30 to their original values. This data reversion is only performed after the two systems reconnect. For now, volume A reverts the semantics of the uncommitted data list to be data that needs to be copied from volume B. For example, blocks 1000 through 2000 need to be copied now from volume B.
12:30 to 13:00	Volume B is primary, volume A is secondary, coupling is non-operational	Volume A does not accept changes because it is a secondary in a nonoperational coupling. volume B is now a primary in a nonoperational coupling, and maintains its own uncommitted data list of the write operations that were performed since it was defined as the primary. For example, if the hosts wrote blocks 1500 through 2500, volume B marks these blocks to be copied to volume A.
13:00	Connectivity resumes	Volume B and volume A communicate and volume B merges the lists of uncommitted data. Volume B copies to volume A both the blocks that changed in volume B between 12:30 and 13:00, as well as the blocks that changed in volume A between 12:00 and 12:30. For example, volume B could copy to volume A blocks 1000 through 2500, where blocks 1000 through 1500 would revert to their original values at 12:00 and blocks 1500 through 2500 would have the values written to volume B between 12:30 and 13:00. Changes written to volume A between 12:00 and 12:30 are lost.

- [Reconnection when both sides have the same role](#)
What happens when one side was switched while the link was down?

Reconnection when both sides have the same role

What happens when one side was switched while the link was down?

Situations where both sides are configured to the same role can only occur when one side was switched while the link was down. This is a user error, and the user must follow these guidelines to prevent such a situation:

- Both sides need to change roles before the link is resumed.
- As a safety measure, it is recommended to first switch the primary to secondary.

If the link is resumed and both sides have the same role, the coupling will not become operational.

To solve the problem, the user must use the role switching mechanism on one of the volumes and then activate the coupling.

In this situation, the system behaves as follows:

- If both sides are configured as secondary volumes, a minor error is issued.
- If both sides are configured as primary volumes, a critical error is issued. Both volumes will be updated locally with remote mirroring being nonoperational until the condition is solved.

Remote mirroring

- [Remote mirroring and consistency groups](#)
The consistency group has to be compatible with mirroring.
- [Using remote mirroring for media error recovery](#)
- [Supported configurations](#)
Synchronous mirroring supports the following configurations.

- [I/O performance versus synchronization speed optimization](#)
The synchronization rate can be adjusted, so as to prevent resource exhaustion.
- [Implications regarding volume and snapshot management](#)
Using synchronous mirroring incurs several implications on volume and snapshot management.

Remote mirroring and consistency groups

The consistency group has to be compatible with mirroring.

The following assumptions ensure that consistency group procedures are compatible with the remote mirroring function:

- All volumes in a consistency group are mirrored on the same system (as all primaries on a system are mirrored on the same system).
- All volumes in a consistency group have the same role.
- The last_consistent snapshot is created and deleted system-wide, and therefore, it is consistent across the consistency group.

Note: An administrator can incorrectly switch the roles of some of the volumes in a consistency group, while keeping others in their original role. This is not prevented by the system and is detected at the application level.

Using remote mirroring for media error recovery

If a media error is discovered on one of the volumes of the coupling, the peer volume is then used for a recovery.

Supported configurations

Synchronous mirroring supports the following configurations.

- Either Fibre Channel or iSCSI can serve as the protocol between the primary and secondary volumes. A volume accessed through one protocol can be synchronized using another protocol.
- The remote system must be defined as an XIV® in the remote-target connectivity definitions.
- All the peers of volumes that belong to the same consistency group on a system must reside on a single remote system.
- Primary and secondary volumes must contain the same number of blocks.

I/O performance versus synchronization speed optimization

The synchronization rate can be adjusted, so as to prevent resource exhaustion.

The IBM® XIV® Storage System has two global parameters, controlling the maximum rate used for initial synchronization and for synchronization after nonoperational coupling.

These rates are used to prevent a situation where synchronization uses too many of the system or communication line resources.

This configuration parameter can be changed at any time. These parameters are set by the IBM XIV Storage System technical support representative.

Implications regarding volume and snapshot management

Using synchronous mirroring incurs several implications on volume and snapshot management.

- 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 snapshot.
- Resizing a primary volume 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.
- 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.

Asynchronous remote mirroring

Asynchronous mirroring enables you to attain high availability of critical data through a process that asynchronously replicates data updates that are recorded on 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 relationship. 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.

XIV features both asynchronous mirroring and synchronous mirroring. Asynchronous mirroring is advantageous in various use cases. It represents a compelling mirroring solution 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 to realize better data availability and cost savings.

With synchronous mirroring (first image below), response time (latency) increases as the distance between peers increases, but both peers are synchronized. With asynchronous mirroring (second image below), response time is not sensitive to distance between peers, but the synchronization gap between peers is sensitive to the distance.

Figure 1. Synchronous mirroring extended response time lag





Figure 2. Asynchronous mirroring - no extended response time lag



Note: Synchronous mirroring is covered in [Synchronous remote mirroring](#).

- **Features**
The IBM XIV Storage System blends existing and new XIV technologies to produce an advanced mirroring solution with unique strengths.
- **Technological overview**
The IBM XIV Storage System asynchronous mirroring blends existing and new technologies.
- **Mirroring consistency groups**
Grouping volumes into a consistency group provides a means to maintain a consistent snapshot of the group of volumes at the secondary site.

Features

The IBM® XIV® Storage System blends existing and new XIV technologies to produce an advanced mirroring solution with unique strengths.

The following are highlights of IBM XIV Storage System asynchronous mirroring:

Advanced snapshot-based technology

IBM XIV asynchronous mirroring is based on XIV snapshot technology, which streamlines implementation while minimizing impact on general system performance. The technology leverages functionality that has previously been effectively employed with synchronous mirroring and is designed to support mirroring of complete systems – translating to hundreds or thousands of mirrors.

Mirroring of consistency groups

IBM XIV supports definition of mirrored consistency groups, which is highly advantageous to enterprises, facilitating easy management of replication for all volumes that belong to a single consistency group. This 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. The XIV implementation allows you to combine both approaches because you can define mirrors with a scheduled replication and you can issue manual replication jobs for these mirrors as needed.

Multiple RPOs and multiple schedules

IBM XIV 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

IBM XIV 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

IBM XIV asynchronous mirroring enables the mirroring of volumes and consistency groups that are stores in thin provisioned pools. This applies to both mirroring peers.

Bi-directional mirroring

IBM XIV systems can host multiple mirror sources and targets concurrently, supporting over a thousand mirrors per system. Furthermore, any given IBM XIV Storage System can have mirroring relationships with several other IBM XIV systems. This enables enormous flexibility when setting mirroring configurations.

The number of systems with which the Storage System can have mirroring relationships is specified out side this document (see the IBM XIV Storage System Data Sheet).

Concurrent synchronous and asynchronous mirroring

The IBM XIV Storage System can concurrently run synchronous and asynchronous mirrors.

Easy transition between peer roles

IBM XIV mirror peers can be easily changed between master and slave.

Easy transition from independent volume mirrors into consistency group mirror

The IBM XIV Storage System 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

IBM XIV 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.

- [Asynchronous remote mirroring terminology](#)
- [Specifications](#)

Asynchronous remote mirroring terminology

Mirror coupling (sometimes referred to as *coupling*)

A pairing of storage peers (either volumes or consistency groups) that are engaged in a mirroring relationship.

Master and slave

The roles that correspond with the source and target storage peers for data replication in a mirror coupling. These roles can be changed by a system administrator after a mirror is created to accommodate customer needs and support failover and failback scenarios. A valid mirror can have only one master peer and only one slave peer.

Peer designation

A user-configurable mirroring attribute that describes the designation associated with a coupling peer. The master is designated by default as primary and the slave is designated by default as secondary. These values serve as a reference for the original peer's designation regardless of any role change issued after the mirror is created, but should not be mistaken for the peer's role (which is either master or slave).

Last replicated snapshot

A snapshot that represents the latest state of the master that is confirmed to be replicated to the slave.

Most recent snapshot

A snapshot that represents the latest synchronized state of the master that the coupling can revert to in case of disaster.

Sync job

The mirroring process responsible for replicating any data updates recorded on the master since the last replicated snapshot was taken. These updates are replicated to the slave.

Schedule

An administrative object that specifies how often a sync job is created for an associated mirror coupling.

Interval

A schedule parameter that indicates the duration between successive sync jobs.

RPO

Recovery Point Objective – an objective for the maximal data synchronization gap acceptable between the master and the slave. An indicator for the tolerable data loss (expressed in time units) in the event of failure or unavailability of the master.

RTO

Recovery Time Objective - an objective for the maximal time to restore service after failure or unavailability of the master.

Specifications

The following specifications apply to the mirroring operation:

Minimum link bandwidth

10Mbps.

Recommended link bandwidth

20Mbps and up.

Maximum round trip latency

250ms.

Attaching XIV systems for mirroring

The connection between two XIV systems has to pass via SAN.

Technological overview

The IBM® XIV® Storage System asynchronous mirroring blends existing and new technologies.

The asynchronous mirroring implementation is based on snapshots and features the ability to establish automatic and manual mirroring with the added flexibility to assign each mirror coupling with a different RPO. The ability to specify a different schedule for each mirror independently from the RPO helps accommodate special mirroring prioritization requirements without subjecting all mirrors to the same mirroring parameters. The paragraphs below detail the following IBM XIV asynchronous mirroring aspects, technologies, and concepts:

- The replication scheme
- The snapshot-based technology
- IBM XIV asynchronous mirroring special snapshots
- Initializing IBM XIV asynchronous mirroring
- The mirroring replication unit: the sync job
- Mirroring schedules and intervals
- The manual (ad-hoc) sync job
- Determining mirror state through the RPO
- Mirrored consistency groups
- IBM XIV asynchronous mirroring and pool space depletion
- [Replication scheme](#)
IBM XIV asynchronous mirroring supports establishing mirroring relationships between an IBM XIV Storage System and other XIV systems.
- [Snapshot-based technology](#)
IBM XIV features an innovative snapshot-based technology for asynchronous mirroring that 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.
- [Initializing the mirroring](#)
An XIV mirror is easily created using the CLI or GUI. First, the mirror is created and activated, then an initialization phase starts.
- [The sync job](#)
Data synchronization between the master and slave is achieved through a process run by the master called a sync job.
- [Mirroring schedules and intervals](#)
The IBM XIV Storage System implements a scheduling mechanism that is used to drive a recurring asynchronous mirroring process.
- [The mirror snapshot \(ad-hoc sync job\)](#)
You can manually issue a dedicated command to run a mirror snapshot, in addition to using the schedule-based option.
- [Determining replication and mirror states](#)
The mirror state indicates whether the master is mirrored according to objectives that are specified by the user.
- [Asynchronous mirroring process walkthrough](#)
- [Peers roles](#)
Peers' statuses denote their roles within the coupling definition.
- [Activating the mirroring](#)
The state of the mirroring is derived from the state of its components.


Replication scheme

IBM XIV asynchronous mirroring supports establishing mirroring relationships between an IBM® XIV® Storage System and other XIV systems.

Each of these relationships can be either synchronous or asynchronous and a system can concurrently act as a master in one relationship and act as the slave in another relationship. There are also no practical distance limitations for asynchronous mirroring – mirroring peers can be located in the same metropolitan area or in separate continents.

Each IBM XIV Storage System can have mirroring relationships with other XIV storage systems. Multiple concurrent mirroring relationships are supported with each target.

Figure 1. The replication scheme



Snapshot-based technology

IBM® XIV® features an innovative snapshot-based technology for asynchronous mirroring that facilitates concurrent mirrors with different recovery objectives.

With IBM XIV asynchronous mirroring, write order on the master is not preserved on the slave. As a result, a snapshot taken of the slave 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 master, it is imperative to maintain a consistent replica of the master that can ensure service continuity.

This is achieved through XIV snapshots. XIV asynchronous mirroring employs snapshots to record the state of the master, and calculates the difference between successive snapshots to determine the data that needs be copied from the master to the slave as part of a corresponding replication process. Upon completion of the replication process, a snapshot is taken of the slave and reflects a valid replica of the master.

Below are select technological properties that explain how the snapshot-based technology helps realize effective asynchronous mirroring:

- XIV's support for a practically unlimited number of snapshots facilitates mirroring of complete systems with practically no limitation on the number of mirrored volumes supported
- XIV 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 taken of the master (being either a volume or consistency group), prior to the creation of a new replication process (Sync Job – see below). This snapshot exists only on the master.

last_replicated snapshot (LRS)

This is the most recent snapshot that is confirmed to have been fully replicated to the slave. Both the master and the slave have this snapshot. On the slave, the snapshot is taken upon completion of a replication process, and replaces any previous snapshot with that name. On the master, the *most_recent* snapshot is renamed *last_replicated* after the slave is confirmed to have a corresponding replica of the master's *most_recent* snapshot.

Figure 1. Location of special snapshots



XIV maintains three snapshots per mirror coupling: two on the master and one on the slave. A valid (recoverable) state of the master is captured in the *last_replicated* snapshot, and on an identical snapshot on the slave. The *most_recent* snapshot represents a recent state of the master that needs be replicated next to the slave. The system determines the data to replicate by comparing the master's snapshots.

Initializing the mirroring

An XIV mirror is easily created using the CLI or GUI. First, the mirror is created and activated, then an initialization phase starts.

XIV mirrors are created in standby state and must be explicitly activated. During the Initialization phase, the system generates a valid replica of the state of the master on the slave. Until the Initialization is over, there is no valid replica on the slave to help recover the master (in a case of disaster). Once the Initialization phase ends, a snapshot of the slave is taken. This snapshot represents a valid replica of the master and can be used to restore a consistent state of the master in disaster recovery scenarios.

The Initialization takes the following steps (all part of an atomic operation):

The master start initializing the slave

When a new mirror is defined, a snapshot of the master is taken. This snapshot represents the initial state of the master prior to the issuing of the mirror. The objective of the Initialization is to reflect this state on the slave.


Initialization finishes

Once the Initialization finishes, an ongoing mirroring commences through a sync job.

Acknowledgment

The slave acknowledges the completion of the Initialization to the master.

Figure 1. Asynchronous mirroring over-the-wire initialization



- **Off-line replicating the master onto the slave**

The IBM XIV Storage System allows for this volume replica to be transferred off-line to the slave.

Off-line replicating the master onto the slave

The IBM® XIV® Storage System allows for this volume replica to be transferred off-line to the slave.

At the beginning of the Initialization of the mirror, the user states which volume will be replicated from the master to the slave. This replica of the master is typically much larger than the schedule-based replicas that accumulate differences that are made during a small amount of time. The IBM XIV Storage System allows for this volume replica to be transferred off-line to the slave. This method of transfer is sometimes called "Truck Mode" and is accessible through the *mirror_create* command.

Off-line initialization of the mirror replicates the master onto the slave without being required to utilize the inter-site link. The off-line replication requires:

- Specifying the volume to be mirrored.
- Specifying the initialization type to the mirror creation command.
- Activating the mirroring.

Meeting the above requirements, the system takes a snapshot of the master, and compares it with the slave volume. Only areas where differences are found are replicated as part of the initialization. Then, the slave peer's content is checked against the master and not just automatically considered a valid replica. This check optimizes the initialization time through taking into consideration the available bandwidth between the master and slave and whether the replica is identical to the master volume (that is, there where no writes to the master during the initialization).

The sync job

Data synchronization between the master and slave is achieved through a process run by the master called a sync job.


The sync job updates the slave with any data that was recorded on the master since the latest sync job was created. The process can either be started automatically based on a user-configurable schedule, or manually based on a user-issued command.

When the sync Job is started, a snapshot of the master's state at that time is taken (the *most_recent_snapshot*).

After any outstanding sync job are completed, the system calculates the data differences between this snapshot and the most recent master snapshot that corresponds with a consistent replica on the slave (the `last_replicated_snapshot`). This difference constitutes the data to be replicated next by the sync job.

The replication is very efficient because it only copies data differences between mirroring peers. For example, if only a single block was changed on the master, only a single block will be replicated to the slave.

Figure 1. The asynchronous mirroring sync job



Mirroring schedules and intervals

The IBM® XIV® Storage System implements a scheduling mechanism that is used to drive a recurring asynchronous mirroring process.

Each asynchronous mirror has a specified schedule, and the schedule's interval indicates how often a sync job is created for that mirror.

Asynchronous mirroring has the following features:

- The schedule concept. A *schedule* specifies an interval for automatic creation of sync jobs; a new sync job is normally created at the arrival of a new interval.
- A sync job is not created if another scheduled sync job is running when a new interval arrives
- Custom schedules can be created by users
- Schedule intervals can be set to any of the following values: 30 seconds, 1 min, 2 min, 5 min, 10 min, 15 min, 30 min, 1 hour, 2 hours, 3 hours, 6 hours, 8 hours, 12 hours. The schedule start hour is 00:00.
Note: The IBM XIV Storage System offers a built-in, non-configurable schedule called `min_interval` with a 20s interval. It is only possible to specify a 20s schedule using this predefined schedule.
- When creating a mirror, two schedules are specified - one per peer. The slave's schedule can help streamline failover scenarios - controlled by either XIV or a 3rd party process.
- A single schedule can be referenced by multiple couplings on the same system.
- Sync job creation for mirrors with the same schedule takes place at exactly the same time. This is in contrast with mirrors having different schedules with the same interval. Despite having the same interval, sync jobs for these types of mirrors are not guaranteed to take place at the same time.
- A unique schedule called `Never` is provided to indicate that no sync jobs are automatically created for the pertinent mirror (see below).

Schedules are local to a single XIV system

Schedules are local to the XIV system where they are defined and are set independently of system-to-system relationships. A given source-to-target replication schedule does not mandate an identical schedule defined on the target for reversed replication. To maintain an identical schedule for reverse replication (if the master and slave roles need be changed), independent identical schedules must be defined on both peers.

Schedule sensitivity to timezone difference

The schedules of the peers of a mirroring couple have to be defined in a way that won't be impacted from timezone differences. For example, if the timezone difference between the master and slave sites is two hours, the interval is 3 hours and the schedule on one peer is (12PM, 3AM, 6AM,...), the schedule on the other peer needs to be (2AM, 5AM, 8AM). Although some cases do not call for shifted schedules (for example, a timezone difference of 2 hours and an interval of one hour), this issue can't be overlooked.

The master and the slave also have to have their clocks synchronized (for example using NTP). Avoiding such a synchronization could hamper schedule-related measurements, mainly RPO.

The Never schedule

The system features a special, non-configurable schedule called `Never` that denotes a schedule with no interval. This schedule indicates that no sync jobs are automatically created for the mirror so it is only possible to issue replication for the mirror through a designated manual command.

Note: A manual snapshot mirror can be issued to every mirror that is assigned a user-defined schedule.

The mirror snapshot (ad-hoc sync job)

You can manually issue a dedicated command to run a mirror snapshot, in addition to using the schedule-based option.

This type of mirror snapshot can be issued for a coupling regardless of whether it has a schedule. The command creates a new snapshot on the master and manually initiates a sync job that is queued behind outstanding sync jobs.

The mirror snapshot:

- Accommodates a need for adding manual replication points to a scheduled replication process.
- Creates application-consistent replicas (in cases where consistency is not achieved via the scheduled replication).

The following characteristics apply to the manual initiation of the asynchronous mirroring process:

- Multiple mirror snapshot commands can be issued – there is no maximum limit on the number of mirror snapshots that can be issued manually.
- A mirror snapshot running when a new interval arrives delays the start of the next interval-based mirror scheduled to run, but does not cancel the creation of this sync job.
 - The interval-based mirror snapshot will be canceled only if the running snapshot mirror (ad-hoc) has never finished.

Other than these differences, the manually initiated sync job is identical to a regular interval-based sync job.

Determining replication and mirror states

The mirror state indicates whether the master is mirrored according to objectives that are specified by the user.

As asynchronous mirroring endures a gap that may exist between the master and slave states, the user must specify a restriction objective for the mirror – the RPO – or Recovery Point Objective. The system determines the mirror state by examining if the master's replica on the slave. The mirror state is considered to be OK only if the master replica on the slave is newer than the objective that is specified by the RPO.

- [RPO and RTO](#)
The evaluation of the synchronization status is done based on the mirror's RPO value. Note the difference between RPO and RTO.
- [Mirror status values](#)
The mirror status is determined based on the mirror state and the mirroring status.
- [Mirrored consistency groups](#)
IBM XIV enables mirrored consistency groups and mirroring of volumes to facilitate the management of mirror groups.
- [Storage space required for the mirroring](#)
IBM XIV enables to manage the storage required for the mirroring on thin-provisioned pools on both the master and the slave.
- [Pool space depletion](#)
Pool space depletion is a mechanism that takes place whenever the mirroring can no longer be maintained due to lack of space for incoming write requests issued by the host.

RPO and RTO

The evaluation of the synchronization status is done based on the mirror's RPO value. Note the difference between RPO and RTO.

RPO

Stands for Recovery Point Objective and represents a measure of the maximum data loss that is acceptable in the event of a failure or unavailability of the master.

RPO units

Each mirror must be set an RPO by the administrator, expressed in time units. Valid RPO values range between 30 seconds and 24 hours. An RPO of 60 seconds indicates that the slave's state should not be older than the master's state by more than 60 seconds. The system can be instructed to alert the user if the RPO is missed, and the system's internal prioritization process for mirroring is also adjusted.

RTO

Stands for Recovery Target Objective and represents the amount of time it takes the system to recover from a failure or unavailability of the master. The mirror's RTO is not administered in XIV.

Mirror status values

The mirror status is determined based on the mirror state and the mirroring status.

During the progress of a sync job and until it completes, the slave replica is inconsistent because write order on the master is not preserved during replication. Instead of reporting this state as inconsistent, the mirror state is reported based on the timestamp of the slave's last_replicated snapshot as one of the following:

RPO_OK

Synchronization exists and meets its RPO objective.

RPO_Lagging

Synchronization exists but lags behind its RPO objective.

Initializing

Mirror is initializing.

- **Definitions of mirror state and status**

The mirror status is determined based on the mirror state and the mirroring status.

- **Determining mirror status**

The mirroring status denotes the status of the replication process and reflects the activation state and the link state.

Definitions of mirror state and status

The mirror status is determined based on the mirror state and the mirroring status.

Mirror state

During the progress of a sync job and until it completes, the slave replica is inconsistent because write order on the master is not preserved during replication. Instead of reporting this state as inconsistent, the mirror state is reported based on the timestamp of the slave's last_replicated snapshot as one of the following:

Definition of RPO_OK

Synchronization exists and meets its RPO objective.

Definition of RPO_Lagging

Synchronization exists but lags behind its RPO objective.

Initializing

Mirror is initializing.

Mirroring status

The mirroring status denotes the status of the replication process and reflects the activation state and the link state.


Effective recovery currency

Measures as the delta between the current time and the timestamp of the last_replicated_snapshot

Declaring on RPO_OK

The effective recovery currency is positive.

Figure 1. The way RPO_OK is determined
RPO_OK



Declaring on RPO_Lagging

The effective recovery currency is negative.

Figure 2. The way RPO_Lagging is determined
RPO_Lagging



Determining mirror status

The mirroring status denotes the status of the replication process and reflects the activation state and the link state.

The following example portrays the way the mirroring status is determined. The time axis denotes time and the schedule of the sync jobs ($t_0 - t_5$). Both RPO states are displayed in red and green at the upper section of the image.

First sync job - RPO is OK

Time: t_0

A sync job starts. RPO_OK is maintained as long as the sync job ends before t_1 .


Time: t_a

As the sync job ends at t_a , prior to t_1 , the status is RPO_OK.

Effective recovery currency

During the sync job run the value of effective recovery currency (the black graph on the upper section of the image) changes. This value goes up as we are getting farther from t_0 , goes down - to the RPO setting - once the sync job complete, and does not resume climbing as long as the next schedule has arrived.

Figure 1. Determining the asynchronous mirroring status – example part 1



Second sync job - RPO is lagging

Time: t_1

A sync job starts. RPO_OK is maintained as long as the sync job ends before t_2 .

Time: t_2

The sync job should have ended at this point, but it is still running.

The sync job the was scheduled to run on this point in time is cancelled.


Time: t_b

As the sync job ends at t_b , which is after t_2 , the status is RPO_Lagging.

Effective recovery currency

The value of effective recovery currency $k.jpg$ climbing as long as the next sync job hasn't finished.

Figure 2. Determining the asynchronous mirroring status – example part 2



Third sync job - RPO is OK

Time: t_3

A new sync job starts. At this point the status is RPO_Lagging.


Time: t_c

As the sync job ends prior to t_4 , the status is RPO_OK.

Effective recovery currency

The value of effective recovery currency $k.jpg$ climbing until the next sync job has finished (this happens at t_c). This value immediately returns to the RPO setting until the time of the next schedule.

Figure 3. Determining Asynchronous mirroring status – example part 3



The added-value of multiple RPOs

The system bases its internal replication prioritization on the mirror RPO; hence, support for multiple RPOs corresponding with true recovery objectives helps optimize the available bandwidth for replication.

The added-value of multiple Schedules

You can attain a target RPO using multiple schedule interval options. A variable schedule that is decoupled from the RPO helps optimize the replication process to accommodate RPO requirements without necessarily modifying the RPO.

Mirrored consistency groups

IBM® XIV® enables mirrored consistency groups and mirroring of volumes to facilitate the management of mirror groups.

Asynchronous mirroring of consistency groups is accomplished by taking a snapshot group of the master consistency group in the same manner employed for volumes, either based on schedules, or manually through a dedicated command option.

The peer synchronization and status are managed on a consistency group level, rather than on a volume level. This means that administrative operations are carried out on the whole consistency group, rather than on a specific volume within the consistency group. This includes operations such as activation, and mirror-wide settings such as a schedule.

The synchronization status of the consistency group reflects the combined status of all mirrored volumes pertaining to the consistency group. This status is determined by examining the (system-internally-kept) status of each volume in the consistency group. Whenever a replication is complete for all volumes and their state is RPO_OK, the consistency group mirror status is also RPO_OK. On the other hand, if the replication is incomplete or any of the volumes in a mirrored consistency group has the status of RPO_Lagging, the consistency group mirror state is also RPO_Lagging.

Storage space required for the mirroring

IBM® XIV® enables to manage the storage required for the mirroring on thin-provisioned pools on both the master and the slave.

Throughout the course of the mirroring, the last_replicated and most_recent snapshots may exceed the space allocated to the volume and its snapshots. The lack of sufficient space on the master can prevent host writes, where the lack of space on the slave can disrupt the mirroring process itself.

IBM XIV enables to manage the storage required for the mirroring on thin-provisioned pools. This way, The IBM XIV Storage System manages and allocates space according to the schemes described in the [Thin provisioning](#) chapter.

Upon depletion of space on each of the peers, the [Pool space depletion](#) mechanism takes effect.

Pool space depletion

Pool space depletion is a mechanism that takes place whenever the mirroring can no longer be maintained due to lack of space for incoming write requests issued by the host.

Whenever a pool does not have enough free space to accommodate the storage requirements warranted by a new host write, the system runs a multi-step procedure that progressively deletes snapshots within that pool until enough space is made available for a successful completion of the write request.

This multi-step procedure is progressive, meaning that the system proceeds to the next step only if following the execution of the current step, there is still insufficient space to support the write request.

- [Protecting snapshots using deletion priority](#)
Protected snapshots have precedence over other snapshots during the pool space depletion process.
- [Pool space depletion on the master](#)
The depletion procedure on the master takes the following steps.
- [Pool space depletion on the slave](#)
Pool space depletion on the slave means that there is no room available for the last_replicated snapshot. In this case, the mirroring is deactivated.


Protecting snapshots using deletion priority

Protected snapshots have precedence over other snapshots during the pool space depletion process.


The concept of protected snapshots assigns the storage pool with an attribute that is compared with the snapshots' auto-deletion priority attribute. Whenever a snapshot has a deletion priority that is higher than the pool's attribute, it is considered protected.

For example, if the deletion priority of the depleting storage is set to 3, the system will delete snapshots with the deletion priority of 4. Snapshots with priority level 1, 2 and 3 will not be deleted.


Figure 1. The deletion priority of the depleting storage is set to 3



If the deletion priority of the depleting storage is set to 4, the system will deactivate mirroring before deleting any snapshots.



If the deletion priority of the depleting storage is set to 0, the system can delete any snapshot regardless of deletion priority.



Deletion priority conventions

Protecting the deletion priority of the last replicated snapshot

The deletion priority of mirror-related snapshots is set implicitly by the system and cannot be customized by the user (see below).

Last replicated and most recent snapshots

The deletion priority of the asynchronous mirroring last_replicated and most_recent snapshots on the master is set to 1.

Last replicated snapshot on the slave

The deletion priority of the last_replicated snapshot on the slave and the is set to 0 (see below).

Default value of the snapshot protecting CLI

By default, the value of the protected_snapshot_priority parameter of the pool_config_snapshots command is 0.

Changing this value

If the protected_snapshot_priority parameter is changed, the system and user-created snapshots with a deletion priority nominally equal or lower than the protected setting will be deleted only after the internal mirroring snapshots are .

For example, if the protected_snapshot_priority is changed to 1, all system and user-created snapshots with deletion priority 1 (which includes ALL snapshots created by the user, assuming that their deletion priority was not changed) will be protected and will be deleted only after internal mirroring snapshots are.

Other snapshots

Non mirroring-related snapshots are created by default with a deletion priority 1.

Protecting the last replicated snapshots

The last replicated snapshots represent a consistent replica of the master in asynchronous mirroring. Both the master and the slave have a last replicated snapshot, however, these two snapshots are protected differently.

LRS_{slave}

The slave must have an available consistent copy of the master at all times, where the master does not have to have such availability (as the LRS_{master} itself is regarded consistent). As a result, this snapshot is never deleted. Upon pool space depletion on the slave, whenever there is no space for the mirroring process, the pool will be locked.

The deletion priority of the LRS on the slave is 0.

LRS_{master}

The last replicated snapshot on the master is available for deletion during pool space depletion.

The deletion priority of the LRS on the master is 1.

Pool space depletion on the master

The depletion procedure on the master takes the following steps.

Step 1 - deletion of unprotected snapshots

The following snapshots are deleted:

- Regular (not related to mirroring) snapshots
- Snapshots of the mirroring processes that are no longer active
- The snapshot of any snapshot mirror (ad hoc sync job) that has not started yet

The deletion is subject to the deletion priority of the individual snapshot. In the case of deletion priority clash, older snapshots are deleted first.

Success criteria:

The user reattempts operation, re-enables mirroring and resumes replication. If this fails, the system proceeds to step 2 (below).

Step 2 - deletion of the snapshot of any outstanding (pending) scheduled sync job

If replication still does not resume after the actions taken on step 1:

The following snapshots are deleted:

- All snapshots that were not deleted in step 1.

Success criteria:

The system reattempts operation, re-enables mirroring and resumes replication.

Step 3 - automatic deactivation of the mirroring and deletion of the snapshot designated as the mirror most_recent snapshot

If the replication still does not resume:

The following takes place:

- An automatic deactivation of the mirroring
- Deletion of the most_recent snapshot
- An event is generated.

Ongoing ad-hoc sync job

The snapshot created during the ad-hoc sync job is considered as a most_recent snapshot, although it is not named as such and not duplicated with a snapshot in that name. Following the completion of the ad-hoc sync job, and only after this completion, the snapshot is duplicated and the duplicate is named last_replicated.

Upon a manual reactivation of the mirroring process:

1. The mirroring activation state changes to Active
2. A most_recent snapshot is created
3. A new sync job starts

Step 4 - deletion of the last_replicated snapshot

If more space is still required:

The following takes place:

- Deletion of the last_replicated snapshot (on the master)
- An event is generated.

Following the deletion:

1. The mirroring remains deactivated, and must be manually reactivated.
2. The mirroring changes to *change tracking* state. Host I/O to the master are tracked but not replicated
3. The system marks storage areas that were written into since the last_replicated snapshot was created

Step 5 - deletion of the most_recent snapshot that is created when activating the mirroring in *Change Tracking* state

If more space is still required:

The following takes place:

- Deletion of the most_recent snapshot (on the master).
- An event is generated.

Following the deletion:

Deletion of this `most_recent` snapshot in this state leaves the master with neither a snapshot nor a bitmap, mandating full initialization. To minimize the likelihood for such deletion, this snapshot is automatically assigned a special (new) deletion priority level. This deletion priority implies that the system should delete the snapshot only after all other `last_replicated` snapshots in the pool were deleted. Note that the new priority level will only be assigned to a mirror with a consistent Slave replica and not to a mirror that was just created (whose first state is also initialization).

Step 6 - deletion of protected snapshots

If more space is still required:

The following takes place:

- An event is generated.
- Deletion of all protected snapshots, regardless of the mirroring. These snapshots are deleted according to their deletion priority and age.

Following the deletion:

- The master's state changes to **Init** (distinguished from the Initialization phase mirrors start with)
- The system stops marking new writes
- A `most_recent` snapshot is created
- The system creates and runs a sync job encompassing all of the tracked changes tracked
- following the completion of this sync job, a `last_replicated` snapshot is created on the master, and the mirror state changes to `rpo_ok` or `rpo_lagging`, as warranted by the effective RPO

If pool space depletes during the Init:

- The master's state remains Initialization
- An event is generated
- The mirroring is deactivated
- The `most_recent` snapshot is deleted (mandating a Full Initialization)

Upon manual mirroring activation during the Init:

- The master's state remains Initialization
- A `most_recent` snapshot is created
- The system starts a Full Initialization based on the `most_recent` snapshot

Pool space depletion on the slave

Pool space depletion on the slave means that there is no room available for the `last_replicated` snapshot. In this case, the mirroring is deactivated.

Snapshots with a deletion priority of 0 are special snapshots that are created by the system on the slave peer and are not automatically deleted to free space, regardless of the pool space depletion process. The asynchronous mirroring slave peer has one such snapshot: the `last_replicated` snapshot.

Asynchronous mirroring process walkthrough

This section walks you through creating an asynchronous mirroring relationship, starting from the initialization all the way through completing the first scheduled sync job.

- **Step 1**
Time is 01:00 when the command to create a new mirror is issued. In this example, an RPO of 120 minutes and a schedule of 60 minutes are specified for the mirror.
- **Step 2**
After the state of the master is captured, the data that needs to be replicated as part of the Initialization process is calculated. In this example, the master's `most_recent` snapshot represents the data to be replicated through the first sync job (4).
- **Step 3**
During this step, the Initialization Sync Job is well in progress. The master continues to be updated with host writes – the updates are noted in the order they are written – first 1, then 2 and finally 3. The initialization sync job replicates the initial master peer's state to the slave peer (5).
- **Step 4**
Moments later, the initialization sync job completes. After it completes, the slave's state is captured by taking a snapshot: the `last_replicated` snapshot (6). This snapshot reflects the state of the master as captured in the `most_recent` snapshot. In this example, it is the state just before the initialization phase started.
- **Step 5**
During this step, the master's `last_replicated` snapshot is created. The `most_recent` snapshot on the master is renamed `last_replicated` (7) and represents the most recent point-in-time that the master can be restored if needed (because this state is captured in the slave's corresponding snapshot).
- **Step 6**
Based on the mirror's schedule, a new interval arrives in a manner similar to the Initialization phase: host writes are blocked (1), and a new master `most_recent` snapshot is created (2), reflecting the master peer's state at this time.
- **Step 7**
A new sync job is set. The data to be replicated is calculated based on the difference between the master's `most_recent` snapshot and the `last_replicated` snapshot (4).
- **Step 8**
The sync job is in process. During the sync job, the master continues to be updated with host writes (update 5).
- **Step 9**
The sync job is completed. The slave's `last_replicated` snapshot is deleted (6) and replaced (in one atomic operation) by a new `last_replicated` snapshot.


- [Step 10](#)
The sync job is completed with a new last_replicated snapshot representing the updated slave's state (7).
- [Step 11](#)
A new master last_replicated snapshot created. In one transaction - the current last_replicated snapshot on the master is deleted (8) and the most_recent snapshot is renamed the last_replicated (9).

Step 1

Time is 01:00 when the command to create a new mirror is issued. In this example, an RPO of 120 minutes and a schedule of 60 minutes are specified for the mirror.

The mirroring process must first establish a baseline for ensuing replication. This warrants an Initialization process during which the current state of the master is replicated to the slave peer. This begins with the host writes being briefly blocked (1). The state of the master peer can then be captured by taking a snapshot of the master state: the most_recent snapshot (2), which serves as a baseline for ensuing schedule-based mirroring. After this snapshot is created, host writes are no longer blocked and continue to update the storage system (3). At this time, no snapshot exists on the slave yet.


Figure 1. Asynchronous mirroring walkthrough – Part 1



Step 2

After the state of the master is captured, the data that needs to be replicated as part of the Initialization process is calculated. In this example, the master's most_recent snapshot represents the data to be replicated through the first sync job (4).


Figure 1. Asynchronous mirroring walkthrough – Part 2



Step 3

During this step, the Initialization Sync Job is well in progress. The master continues to be updated with host writes – the updates are noted in the order they are written – first 1, then 2 and finally 3. The initialization sync job replicates the initial master peer's state to the slave peer (5).


Figure 1. Asynchronous mirroring walkthrough – Part 3



Step 4

Moments later, the initialization sync job completes. After it completes, the slave's state is captured by taking a snapshot: the last_replicated snapshot (6). This snapshot reflects the state of the master as captured in the most_recent snapshot. In this example, it is the state just before the initialization phase started.

Figure 1. Asynchronous mirroring walkthrough – Part 4




Step 5

During this step, the master's last_replicated snapshot is created. The most_recent snapshot on the master is renamed last_replicated (7) and represents the most recent point-in-time that the master can be restored if needed (because this state is captured in the slave's corresponding snapshot).

When the initialization phase ends, the master and slave peers have an identical restore time point, to which they can be reverted, if needed.

Figure 1. Asynchronous mirroring walkthrough – Part 5




Step 6

Based on the mirror's schedule, a new interval arrives in a manner similar to the Initialization phase: host writes are blocked (1), and a new master most_recent snapshot is created (2), reflecting the master peer's state at this time.

Then, host writes are no longer blocked (3).

The update number (4) occurs after the snapshot is taken and is not reflected in the next sync job. This is shown by the color-shaded cells in the most_recent snapshot figure.

Figure 1. Asynchronous mirroring walkthrough – Part 6



Step 7

A new sync job is set. The data to be replicated is calculated based on the difference between the master's most_recent snapshot and the last_replicated snapshot (4).

Figure 1. Asynchronous mirroring walkthrough – Part 7



Step 8

The sync job is in process. During the sync job, the master continues to be updated with host writes (update 5).

The sync job data is not replicated to the slave in the order by which it was recorded at the master – the order of updates on the slave is different.


Figure 1. Asynchronous mirroring walkthrough – Part 8



Step 9

The sync job is completed. The slave's last_replicated snapshot is deleted (6) and replaced (in one atomic operation) by a new last_replicated snapshot.

Figure 1. Asynchronous mirroring walkthrough – Part 9




Step 10

The sync job is completed with a new last_replicated snapshot representing the updated slave's state (7).

The slave's last_replicated snapshot reflects the master's state as captured in the most_recent snapshot. In this example, it is the state at the beginning of the mirror schedule's interval.

Figure 1. Asynchronous mirroring walkthrough – Part 10




Step 11

A new master last_replicated snapshot created. In one transaction - the current last_replicated snapshot on the master is deleted (8) and the most_recent snapshot is renamed the last_replicated (9).

The interval sync process is now complete - the master and slave both have an identical restore time point to which they can be reverted if needed.

Figure 1. Asynchronous mirroring walkthrough – Part 11



Peers roles

Peers' statuses denote their roles within the coupling definition.

After creation, a coupling has exactly one peer that is set to be the master peer, and exactly one peer that is set to be the slave peer. Each of the peers can have the following available statuses:

None

The peer is not part of a coupling definition.

Master

The actual source peer in a replication coupling. This type of peer serves host requests, and is the source for synchronization updates to the slave. A master peer can be changed to a slave directly while in asynchronous mirroring.

Slave

The actual target peer in a replication coupling. This type of peer does not serve host requests, and accepts synchronization updates from a corresponding master. A slave can be changed to a master directly while in asynchronous mirroring.

Activating the mirroring

The state of the mirroring is derived from the state of its components.

The remote mirroring process hierarchically manages the states of the entities that participate in the process. It manages the states for the mirroring based on the states of the following components:

- [Link](#)
- [Activation](#)

The following mirroring states are possible:

Non-operational

The coupling state is defined as non-operational if at least one of the following conditions is met:

- The activation state is standby.
- The link state is error.
- The slave peer is locked.

Operational

All of the following conditions must be met for the coupling to be defined as operational:

- The activation state is active.
 - The link is OK.
 - The peers have different roles.
 - The slave peer is not locked.
- [Link states](#)
The link state is one of the factors determining the coupling operational status.
 - [Activation states](#)
When the coupling is created, its activation is in standby state. When the coupling is enabled, its activation is in active state.
 - [Deactivating the coupling](#)
Deactivating the coupling stops the mirroring process.

Link states

The link state is one of the factors determining the coupling operational status.

The *link state* reflects the connection from the master to the slave. A failed link or a failed slave system both manifest as a link error. The link state is one of the factors determining the coupling operational status.

The available link states are:

OK

The link is up and functioning.

Error

The link is down.

Activation states

When the coupling is created, its activation is in standby state. When the coupling is enabled, its activation is in active state.

Standby

When the coupling is created, its activation is in standby state.

The synchronization is disabled:

- Sync jobs do not run.
- No data is copied.
- The coupling can be deleted.

Active

The synchronization is enabled:

- Sync jobs can be run.
- Data can be copied between peers.

Regardless of the activation state:

- The mirroring type can be changed to synchronous.
- Peer roles can change.

Deactivating the coupling

Deactivating the coupling stops the mirroring process.

The mirroring is terminated by deactivating the coupling, causing the system to:

- Terminate, or delete the mirroring
- Stop the mirroring process as a result of:
 - A planned network outage
 - An application to reduce network bandwidth
 - A planned recovery test

The deactivation pauses a running sync job and no new sync jobs will be created as long as the active state of the mirroring is not restored. However, the deactivation does not cancel the interval-based status check by the master and the slave. The synchronization status of the deactivated coupling is calculated on the start of each interval, as if the coupling was active.

Deactivating a coupling while a sync job is running, and not changing that state before the next interval begins, leads to the synchronization status becoming RPO_Lagging, as described in the following outline. Upon the deactivation:

On the master

The activation state changes to standby; replication pauses (and records where it paused); replication resumes upon activation.

Note: An ongoing sync job resumes upon activation, no new sync job will be created until the next interval.

On the slave

Not available.

Regardless of the state of the coupling:

- Peer roles can be changed

Note: For consistency group mirroring: deactivation pauses all running sync jobs pertaining to the consistency group. It is impossible to deactivate a single volume sync job within a consistency group.

Mirroring consistency groups

Grouping volumes into a consistency group provides a means to maintain a consistent snapshot of the group of volumes at the secondary site.

The following assumptions make sure that consistency group semantics work with remote mirroring:

Consistency group-level management

Mirroring of consistency groups is managed on a consistency group level, rather than on a volume level. For example, the synchronization status of the consistency group is determined after examining all mirrored volumes that pertain to the consistency group.

Starting with an empty consistency group

Only an empty consistency group can be defined as a mirrored consistency group. If you want to define an existing non-empty consistency group as mirrored, the volumes within the consistency group must first be removed from the consistency group and added back only after the consistency group is defined as mirrored.

Adding a volume to an already consistency group

Only mirrored volumes can be added into a mirrored consistency group. This operations requires the following:

- Volume peer is on the same system as the peers of the consistency group
- Volume replication type is identical to the type used by the consistency group. For example, async_interval.
- Volume belongs to the same storage pool of the consistency group
- Volume has the same schedule as the consistency group
- Volume has the same RPO as the consistency group
- Volume and consistency group are in the same synchronization status (SYNC_BEST_EFFORT for synchronous mirroring, RPO OK for asynchronous mirroring)

If the mirrored consistency group is configured with a user-defined schedule, meaning not using the Never schedule:

Mirrored consistency group or volume should not have non-started snapshot mirrors, non-finished snapshot mirrors (ad hoc sync jobs), or both.
If the mirrored consistency group is configured with a Never schedule:
Mirrored consistency group or volume should not have non-started, non-finished snapshot mirrors, non-finished snapshot mirrors (ad hoc sync jobs), or both. The status of the mirrored consistency group shall be Initialization until the next sync job is completed.

Adding a mirrored volume to a non-mirrored consistency group

It is possible to add a mirrored volume to a non-mirrored consistency group, and it will retain its mirroring settings.

A single sync job for the entire consistency group

The mirrored consistency group has a single sync job for all pertinent mirrored volumes within the consistency group.

Location of the mirrored consistency group

All mirrored volumes in a consistency group are mirrored on the same system.

Retaining mirroring attributes of a volume upon removing it from a mirrored consistency group

When removing a volume from a mirrored consistency group, the corresponding peer volume is removed from the peer consistency group. Mirroring is retained (same configuration as the consistency group from which it was removed). Peer volume is also removed from peer consistency group. Ongoing consistency group sync jobs will continue.

Mirroring activation of a consistency group

Activation and deactivation of a consistency group affects all consistency group volumes.

Role updates

Role updates concerning a consistency group affects all consistency group volumes.

Dependency of the volume on its consistency group

- It is not possible to directly activate, deactivate, or update role of a given volume within a consistency group from the UI.
- It is not possible to directly change the interval of a given volume within a consistency group.
- It is not possible to set independent mirroring of a given volume within a consistency group.

Protecting the mirrored consistency group

Consistency group-related commands, such as moving a consistency group, deleting a consistency group and so on, are not allowed as long as the consistency group is mirrored. You must remove mirroring before you can delete a consistency group, even if it is empty.

- [Setting a consistency group to be mirrored](#)
Volumes added to a mirrored consistency group have to meet some prerequisites.
- [Creating a mirrored consistency group](#)
The process of creating a mirrored consistency group comprises the following steps.
- [Adding a mirrored volume to a mirrored consistency group](#)
After the volume is mirrored and shares the same attributes as the consistency group, you can add the volume to the consistency group after certain conditions are met.
- [Removing a volume from a mirrored consistency group](#)
Removal of a volume from a mirrored consistency group is easy and preserves volume mirroring.

Setting a consistency group to be mirrored

Volumes added to a mirrored consistency group have to meet some prerequisites.

Volumes that are mirrored together as part of the same consistency group share the same attributes:

- Target
- Pool
- Sync type
- Mirror role
- Schedule
- Mirror state
- Last_replicated snapshot timestamp

In addition, their snapshots are all part of the same last_replicated snapshot group.

To obtain the consistency of these attributes, setting the consistency group to be mirrored is done by first creating a consistency group, then setting it to be mirrored and only then populating it with volumes. These settings mean that adding a new volume to a mirrored consistency group requires having the volume set to be mirrored exactly as the other volumes within this consistency group, including the last_replicated snapshot timestamp (which entails an RPO_OK status for this volume).

Note: A non-mirrored volume cannot be added to a mirrored consistency group. It is possible, however, to add a mirrored volume to a non-mirrored consistency group, and have this volume retain its mirroring settings.

Creating a mirrored consistency group

The process of creating a mirrored consistency group comprises the following steps.

Step 1

Define a consistency group as mirrored (the consistency group must be empty).

Step 2

Activate the mirror.

Step 3

Add a corresponding mirrored volume into the mirrored consistency group. The mirrored consistency group and the mirrored volume must have the following identical parameters:

- Source and target
- Pools
- Mirroring type
- RPO
- Schedule names (both local and remote)
- Mirror state is RPO_OK
- Mirroring status is Activated

Note: It is possible to add a mirrored volume to a non-mirrored consistency group. In this case, the volume retains its mirroring settings.

Adding a mirrored volume to a mirrored consistency group

After the volume is mirrored and shares the same attributes as the consistency group, you can add the volume to the consistency group after certain conditions are met.

The following conditions must be met:

- The volume is on the same system as the consistency group
- The volume belongs to the same storage pool as the consistency group
- Both the volume and the consistency group do not have outstanding sync jobs, either scheduled or manual (ad hoc)
- The volume and consistency group have the same synchronization status (synchronized="best_effort" and async_interval="rpo_ok")
- The volume's and consistency group's special snapshots, most_recent and last_replicated, have identical timestamps (this is achieved by assigning the volume to the schedule that is utilized by the consistency group)
- In the case that the consistency group is assigned with schedule="never", the status of the consistency group is initialization as long as no sync job has run.

Removing a volume from a mirrored consistency group

Removal of a volume from a mirrored consistency group is easy and preserves volume mirroring.

When you remove a volume from a mirrored consistency group, the corresponding peer volume is removed from the peer consistency group; mirroring is retained with the same configuration as the consistency group from which it was removed. All ongoing consistency group's sync jobs keep running.

Multi-site mirroring

Multi-site mirroring is an IBM® XIV® Storage System technology that allows customers to set High Availability and Disaster Recovery solutions over multiple sites, keeping 3 copies of their data.

Key features of multi-site mirroring are:

Concurrent multiple multi-site mirroring

- The IBM XIV approach to multi-site mirroring includes 3 peers with one synchronous and two asynchronous replications among them (one of them at standby).
- Multiple multi-site mirroring configurations run concurrently per system, each with separate mirror peers.
- The source runs 2 concurrent mirrors into 2 different destinations
- Any given system can be represented in several multi-site configurations, each referencing different systems.
- A system can host mirroring peers with different roles in different multi-site configurations.

Extensibility

- Any existing two-way mirroring relation (synchronous or asynchronous) can be extended to three-way mirroring, with no need to disrupt the existing mirror relation.
Note: Three-way mirroring cannot be configured on a mirrored consistency group, but can be configured on a local consistency group.
- The multi-site mirroring relation is created based on an already existing target connectivity.

Maintainability

- If one mirror of the multi-site mirror fails, the other mirror continues.
- [Multi-site mirroring terminology](#)
The multi-site mirroring technology introduces some new terms, in addition to those mentioned in the synchronous and asynchronous mirroring chapters.
- [Multi-site mirroring technological overview](#)
The IBM XIV enables replication of a volume to two peer volumes that reside on other systems.

Multi-site mirroring terminology

The multi-site mirroring technology introduces some new terms, in addition to those mentioned in the synchronous and asynchronous mirroring chapters.

Master (Source)

The volume that is mirrored.

Substitute master (Secondary source)

The volume that synchronously mirrors the source.

Slave (Destination)

The volume that asynchronously mirrors the source.

Terminology of synchronous and asynchronous mirroring

Some of the concepts discussed on this chapter were introduced in previous mirroring chapters. For a summary of the terminology these chapters use, see here:

- [Remote mirroring basic concepts.](#)
- [Asynchronous remote mirroring terminology.](#)


Multi-site mirroring technological overview

The IBM® XIV® enables replication of a volume to two peer volumes that reside on other systems.

Components hierarchy

The multi-site relation clearly defines each system and the role it plays during a disaster recovery scenario.

Figure 1. The hierarchy of multi-site mirroring components



The substitute master (System B) is synchronously mirrored with the master (System A), and takes on the role of the master to the slave system (System C) when the master (System A) becomes unavailable.

C is a replica of either A or B. If the A-C mirroring relation is active, then the B-C mirroring is inactive, and vice versa. That is, both A and B cannot write to C at the same time.

A-B mirror

The synchronous mirroring relation between the master and substitute master.

A-C mirror

The asynchronous mirroring relation between the master and slave.

B-C mirror

The asynchronous mirroring relation between the substitute master and the slave.

This mirroring relation is also known as the *Standby* mirror.

The mirroring relation that comprises the B-C mirroring relation can be either of the following types:

- Standby mirror - the third mirror of the multi-site mirroring definition, which is defined in advance
- Live mirror - an operational mirroring relation, which becomes operational only by request in case of disaster recovery

Defining the standby mirroring relation in advance requires that the target connectivity between B and C (or at least its definitions) needs to be in place between all systems when the multi-site mirroring relation is configured.

[Table 1](#) and [Figure 1](#) display the roles of each of the systems that participate in the multi-site mirroring relations.

Table 1. The mirroring relations that comprise the multi-site mirroring

System	Role	A-B	A-C	B-C
A	Source	Synchronous mirroring relation. System A is the <i>master</i> . The mirror is <i>active</i> .	Asynchronous mirroring relation. System A is the <i>master</i> . The mirror is <i>active</i> .	
B	Secondary source	Synchronous mirroring relation. System B is the <i>slave</i> . The mirror is <i>active</i> .		Asynchronous mirroring relation. System B is the <i>master</i> . The mirror is <i>standby</i> .
C	Destination		Asynchronous mirroring relation. System C is the <i>slave</i> . The mirror is <i>active</i> .	Asynchronous mirroring relation. System C is the <i>slave</i> . The mirror is <i>standby</i> .

Note: Currently, multi-site mirroring does not support consistency group mirroring. However, it can be supported when volume mirroring is used and the volume is a part of a local consistency group.

Multi-site mirroring states

The IBM XIV multi-site mirroring technology has multiple states, or conditions, of operation. While each individual mirroring definition has its own state, the multi-site mirroring definition has a global state, too.

Global states

The following states are applicable to the global states of the mirrors:

Init

All mirroring definitions are ready to start transferring data.

Operational

A steady state where both A-B and A-C are *Active*.

Degraded

If both A-B and A-C are active and synchronized, but A-C is *RPO lagging*, the mirroring state is *degraded*.

Inactive

When both A-B and A-C are *inactive*, the mirroring state is *inactive*.

Compromised

These are possible reason for a *compromised* state:

Disconnection

The link is *down* for either A-B or A-C.

Resync

Either A-B or A-C are in *resync* and the substitute master did not yet take ownership.

Following a partial change of role

There was a role change on either A-B or A-C.

Substitute master and slave states

The following states are applicable to substitute master and slave states:

Connected

The mirror with the master system is in a *connected* state.

Disconnected

The mirror with the master system is in a *disconnected* state.

Standby mirroring states

The following states are applicable to the standby mirror:

Up

The standby mirror is defined and connected.

Down

The standby mirror is defined and disconnected.

NA

The standby mirror is not defined.

IBM Hyper-Scale Mobility

IBM Hyper-Scale Mobility enables a 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. Consider 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.
- [The IBM Hyper-Scale Mobility process](#)
This section walks you through the IBM Hyper-Scale Mobility process.

The IBM Hyper-Scale Mobility process

This section walks you through the IBM Hyper-Scale Mobility process.

Hyper-Scale Mobility moves a volume from one system to another, while the host is 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 depicted in [Figure 1](#) and explained in detail in [Table 1](#).

For an in-depth practical guide to using IBM Hyper-Scale Mobility, see the Redbooks publication [IBM Hyper-Scale Mobility Overview and Usage](#).

Figure 1. Flow of the IBM Hyper-Scale Mobility




Table 1. The IBM Hyper-Scale Mobility process

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 Initializing state of synchronous mirroring (see Synchronous mirroring statuses). 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 IBM Hyper-Scale Mobility 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 wholly copied to the destination. This state is similar to the Synchronized state of synchronous mirroring (see Synchronous mirroring statuses).
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.	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 relationship.	

Data-at-rest encryption

The IBM® XIV® Storage System utilizes full disk encryption for regulation compliance and security audit readiness.

Data-at-rest encryption protects against the potential exposure of XIV system sensitive data on discarded or stolen media. The encryption ensures that the data cannot be read, as long as its encryption key is secured. This feature complements physical security at the customer site, protecting the customer from unauthorized access to the data.

The encryption of the disk drives is transparent to hosts that are attached to the IBM XIV Storage System, and does not affect either their management or performance. The term data-in-flight refers to I/Os that are anywhere between the network interfaces, memory and Infiniband backbone. This type of data is not encrypted.

Common use cases that prompt the protection of data-at-rest are:

- Unauthorized access at Service providers SAN's with consolidated storage (e.g. disk theft)
- Component rotation:
 - Protect data following its physical removal from an IBM XIV Storage System at a customer site
 - Prevent discarded media from being compromised (e.g. failed disk)
- New component add - upgrade (MES) of SSD or module should maintain the encryption capabilities of IBM XIV Storage System.
- [HIPAA compatibility](#)
IBM XIV Storage System complies with the following security requirements and standards.

HIPAA compatibility

IBM® XIV® Storage System complies with the following security requirements and standards.

The IBM XIV Storage System Data-at-Rest encryption complies with HIPAA Federal requirements as follows:

- User data is inaccessible without XIV system specific keying material.
- Physical separation of encryption keys from encrypted data, by using an external key server
- Cryptographic keys may be replaced at the user's initiative
- All keys stored must be wrapped and stored in ciphertext (not reside in plain text or hidden/obfuscated)
- AES 256 encryption is used to wrap keys and encrypt data, RSA 2048 encryption is used for public key cryptography
- Encryption configuration and settings must be auditable, thus the related information and notifications should be kept in events log

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 XIV Management Tools**, which includes **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 graphical user interface (GUI) for managing and monitoring multiple storage systems in real time.
- **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.

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 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. Each of the following destinations must be defined as described:

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 information](#)
- [Event notification rules](#)
- [Event notification gateways](#)

Event information

Events are created by various processes, including the following:

- 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

Nonurgent attention is required to verify that there is no problem

Informational

Normal working procedure event

The IBM® XIV® 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 IBM® XIV® 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 an XCLI command or the IBM XIV Storage Management GUI. 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.

Event notification gateways

Event notifications can be sent by SMS, e-mail, or SNMP manager. This step defines the gateways that will be used to send e-mail or SMS.

E-mail (SMTP) gateways

Several e-mail gateways can be defined to enable notification of events by e-mail. By default, the IBM® XIV® Storage System attempts to send each e-mail 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 e-mail 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 e-mail mailbox.

E-mail-to-SMS gateways

SMS messages can be sent to cell phones through one of a list of e-mail-to-SMS gateways. One or more gateways can be defined for each SMS destination.

Each such e-mail-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 e-mail message.

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. Available user roles

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
Storage administrator	The storage administrator has permission to all functions, except: • Maintenance of physical components or changing the status of physical components • Only the predefined administrator, named <i>admin</i>, can change the passwords of other users	Storage administrators are responsible for all administration functions.
Technician	The technician is limited to the following tasks: • Physical system maintenance • Phasing components in or out of service	Technicians maintain the physical components of the system. Only one predefined technician is specified per system. Technicians are IBM® XIV® Storage System technical support team members.

Notes:

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](#)
- [Predefined users](#)
- [User information](#)

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.

Predefined users

There are several predefined users configured on the IBM® XIV® Storage System. These users cannot be deleted.

Storage administrator

This user id provides the highest level of customer access to the system.

Predefined user name: admin

Default password: adminadmin

Technician

This user id is used only by IBM XIV Storage System service personnel.

Predefined user name: technician

Default password: Password is predefined and is used only by the IBM XIV Storage System technicians.

Note: Predefined users are always authenticated by the IBM XIV 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. for explanations of each role.

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: ~!@#\$\$%^&*()_+ -={}|;:

<>? ,./\[] .

E-mail

E-mail is used to notify specific users about events through e-mail messages. E-mail addresses must follow standard addressing procedures. E-mail is optional. Range: Any legal e-mail address.

Phone and area code

Phone numbers are used to send SMS messages to notify specific users about events. Phone numbers and area codes can be a maximum of 63 digits, hyphens (-) and periods (.) Range: Any legal telephone number; The default is N/A

User authentication and access control

IBM® XIV® 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 IBM XIV Storage System, regardless of the authentication mode.

- [Native authentication](#)
Native authentication is the default mode for authenticating users and user groups.
- [LDAP authentication](#)
Lightweight Directory Access Protocol (LDAP) support enables the IBM XIV Storage System to authenticate users through an LDAP repository.
- [LDAP authentication logic](#)
The LDAP authentication process consists of several key steps.

Native authentication

Native authentication is 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 IBM® XIV® Storage System to authenticate users through an LDAP repository.

When LDAP authentication is enabled, the username and password of a user accessing the IBM XIV Storage System (through CLI or GUI) are used by the IBM XIV system to login into a specified LDAP directory. Upon a successful login, the IBM XIV Storage System retrieves the user's IBM XIV group membership data stored in the LDAP directory, and uses that information to associate the user with an IBM XIV administrative role.

The IBM XIV group membership data is stored in a customer defined, pre-configured attribute on the LDAP directory. This attribute contains string values which are associated with IBM XIV administrative roles. These values might be LDAP Group Names, but this is not required by the IBM XIV Storage System. The values the attribute contains, and their association with IBM XIV administrative roles, are also defined by the customer.

Supported domains

The IBM XIV Storage System supports LDAP authentication of the following directories:

- 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 IBM XIV Storage System allows for more than one role for the Storage Administrator and the Read-Only roles.

The predefined XIV administrative IDs “admin” and “technician” are always authenticated by the IBM XIV storage system, whether or not LDAP authentication is enabled.

Responsibilities division between the LDAP directory and the storage system

Following are responsibilities and data maintained by the IBM XIV system and the LDAP directory:

LDAP directory

- Responsibilities - user authentication for IBM XIV users, and assignment of IBM XIV related group in LDAP.
- Maintains - Users, username, password, designated IBM XIV related LDAP groups associated with IBM XIV Storage System.

IBM XIV Storage System

- Responsibilities - Determination of appropriate user role by mapping LDAP group to an IBM XIV role, and enforcement of IBM XIV user system access.
- Maintains - Mapping of LDAP group to IBM XIV role.


LDAP authentication logic

The LDAP authentication process 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



User validation

During the login, the system validates the 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, base 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.

The main idea of multi-tenancy is to allow the allocation of storage resources to several independent administrators with the assurance that one administrator cannot access resources associated with another administrator.

This resource allocation is best described as 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, pools, hosts/clusters, targets, etc. may be associated. The domain restricts the resources a user can manage to those associated with the domain.

A domain maintains the user relationships that exist at the storage system level, as shown in the following figure.



A *domain administrator* is a user who is associated with a domain. The domain administrator is restricted to performing operations on objects associated with a specific domain.

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.

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. They can access the same objects under the same restrictions.

Global administrator

The global administrator is 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

User that is not associated with any domain

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.

The domain administrator

The domain administrator has the following attributes:

- Prior to its association with a domain, the future domain administrator (now 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 admin'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

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's pools
 - The sum of the hard capacity of the system's domains cannot exceed the XIV® system's total hard capacity
 - The sum of the soft capacity of the system's domains cannot exceed the XIV system's total soft capacity
- *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 in 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.

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 the 'Platform and application integration' section .

- [VMware Virtual Volumes](#)

XIV is now ready for *VMware Virtual Volumes (VVOs)*. VMware Virtual Volumes (VVOs) is a feature of VMware vSphere, based on a new storage architecture, that associates a single VM with multiple LUNs.

- [Integration with Microsoft Azure Site Recovery](#)

Microsoft Azure Site Recovery (ASR) solution helps you protect important applications by coordinating the replication and recovery of private clouds across sites.

VMware Virtual Volumes

XIV is now ready for *VMware Virtual Volumes (VVols)*. VMware Virtual Volumes (VVols) is a feature of VMware vSphere, based on a new storage architecture, that associates a single VM with multiple LUNs.

With VVols, the VMware vCenter (Web Client) administrator can offload VM-granular snapshots and cloning to IBM Storage, automate IBM storage provisioning by workload-aware policy, and apply VM-granular backup and *in-place* restore based on IBM Storage snapshots. IBM Storage administrators can define and publish workload-specific storage services to vCenter, and scale down management efforts, enjoying fully automatic volume life cycle management. Lastly, VVols is "elastic", meaning that Storage administrators do not need to pre-allocate large capacity for datastores. Instead, storage is instantly and automatically allocated (and reclaimed) on demand, at exactly the right amount.

VMware VVols automation is based on *VMware vSphere APIs for Storage Awareness (VASA)*. The *IBM Storage Provider for VMware VASA* is a feature of the *IBM Storage Integration Server*, and will support the orchestration of all VVols operations with XIV. For more information on the IBM Storage Provider for VMware VASA and the IBM Storage Integration Server, refer to the *IBM Storage Provider for VMware VASA* (http://www-01.ibm.com/support/knowledgecenter/STJTAG/hsg/hsg_vasa_kcwelcome.html) and *IBM Storage Integration Server* (http://www-01.ibm.com/support/knowledgecenter/STJTAG/hsg/hsg_isis_kcwelcome.html) documentation.

For a preview of *VMware Virtual Volumes (VVols)*, see <http://blogs.vmware.com/vsphere/2012/10/virtual-volumes-vvols-tech-preview-with-video.html>.

Prerequisites for working with VVols

Upon availability, a VVols deployment will require VVols-capable storage arrays and a VASA Provider.

- Make sure the following software and server versions are installed:
 - XIV version 11.5.1, and later
 - IBM Storage Integration Server version 2.0, and up (VASA 2.0-compliant)
 - VMware vCenter and VMware ESX servers
 - VMware vSphere Client
- Deployment of an IBM Storage provider for VASA (incorporated in the IBM Storage Integration Server)
- Define a Storage Integration Administrator (SIA) user role.

Integration with Microsoft Azure Site Recovery


Microsoft Azure Site Recovery (ASR) solution helps you protect important applications by coordinating the replication and recovery of private clouds across sites.

IBM XIV Storage System v11.6.0 supports Microsoft Azure Site Recovery, enabling customers using Microsoft System Center Virtual Machine Manager (SCVMM) to seamlessly orchestrate and manage XIV replication and disaster recovery. Support for Microsoft Azure Site Recovery is based on XIV support for SMI-S v1.6 (<http://www.snia.org/ctf/conformingproviders/ibm.html#sftw4>).

The SCVMM ASR integrates with storage solutions, such as IBM® XIV® CIM Agent, to provide site-to-site disaster recovery for Hyper-V environments by leveraging the SAN replication capabilities that are natively offered by IBM XIV storage systems. It orchestrates replication and failover for virtual machines managed by SCVMM.

SCVMM ASR uses the IBM XIV Remote Mirroring feature through SMI-S to create and manage the replication groups. IBM XIV Remote Mirroring is a host-independent, array-based data mirroring solution that enables affordable data distribution and disaster recovery for applications. With this feature, the users can copy virtual volumes from one IBM XIV storage system to another.

Figure 1. Overview of Microsoft Azure Site Recovery support



Software upgrade

Non disruptive code load (hot upgrade) enables the IBM® XIV® Storage System to upgrade its software from a current version to a newer version without disrupting application service.

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 damaged. The upgrade requires that neither data migration nor rebuild processes are run, and that all internal network paths are active.

During the non disruptive code load process there is a point in time dubbed the 'upgrade-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.

Following are notable characteristics of the 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 online to the application/host.

The duration of the upgrade process is affected by the following factors:

- The upgrade process requires that you reduce all I/Os. If there are a lot of I/Os in the system, or there are slow disks, the system might not be able to stop the I/Os fast enough, so it will restart them and try again after a short while, taking into consideration some retries.
- The upgrade process installs a valid version of the software and then retains its local configuration. This process might take a considerable amount of time, depending on the future changes in the structure of the configuration.

Prerequisites and constraints

- The process cannot run if a data migration process or a rebuild process is active. An attempt to start the upgrade process when either a data migration or a rebuild process is active will fail.
- Generally, everything that happens after the point-of-no-return is treated as if it happened after the upgrade is over.
- 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.
- Prior to the point-of-no-return, a manual abort of the upgrade is available.

Effect on mirroring

Mirrors are automatically deactivated 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 will be aborted automatically.

Handling module or disk failure during the upgrade

If the failure occurs before the point-of-no-return, it will abort the upgrade. If it happens after that point, the failure is treated as if it happened after the upgrade is over.

Handling power failure during the upgrade

As for power failure before the point-of-no-return - power is being monitored during the time the system prepares for the upgrade (before the point-of-no-return). If a power failure is detected, the upgrade will be aborted and the power failure will be taken care of by the old version.

- **Preparing for upgrade**

The IBM XIV Storage System upgrades the system code without disconnecting active hosts or stopping I/O operations.

Preparing for upgrade

The IBM® XIV® Storage System upgrades the system code without disconnecting active hosts or stopping I/O operations.

Important: The upgrade must be performed only by an authorized IBM service technician.

Preparing for the upgrade

Before the code load (upgrade), fulfill the following prerequisites by verifying that:

1. The multipathing feature (provided by the operating system) is working on the host.
2. There are paths from the host to at least two different interface modules on IBM XIV.
3. There is no more than a single initiator in each zone (SAN Volume Controller attached to IBM XIV is an exception).
4. The host was attached to IBM XIV using the xiv_attach utility.
 - This is mandatory.
 - This applies to both installable HAK and portable HAK.
 - Exceptions to this prerequisite are supported platforms, for which no HAK is available (for example, VMware or Linux on Power systems).
5. The minimal version of the "IBM XIV Host Attachment Kit for Windows" is 1.5.3. This version prevents Windows hosts from a potential loss of access.
6. In a case of IBM XIV uses 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.
7. 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 IBM XIV Storage System is not supported.

Be aware of the following:

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. The mirroring is automatically suspended and resumed for a short while during the upgrade.
4. Mirroring from 10.2.4.x to 10.2.1.x or older versions is not supported.
5. There are special considerations where MS Geo Cluster is involved. Contact IBM support for more details.

Recommended practices:

1. It is highly recommended to have the latest XIV Host Attachment Kit installed on the host. Each release of a HAK fixes known issues on older versions.
 - Being on an older level means being exposed to problems already found and fixed.
2. It is recommended to follow the zoning best practices of IBM XIV as described in the Redbook and in the XIV Host Attachment Guide.
3. Best practice is to follow the OS Provider recommendations regarding service packs and storage-related Hot Fixes. These fixes are released from time to time by the OS provider and are outside of IBM XIV control. Some fixes are listed on the release notes of the latest available HAK and have to be applied before the upgrade.
4. It is recommended to keep your host system with up-to-date BIOS and HBA drivers.
5. It is recommended to perform the upgrade on time during which the workload is relatively low.

Availability of TA support:

1. If entitled with a TA service for your IBM XIV Storage System, contact your assigned Technical Advisor when planning for a code upgrade.

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.

- **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.



Work-embedded e-Learning

Learning and tutorials

These e-Learning modules introduce you to the IBM® XIV® Storage System product, define key concepts, and provide an overview of related management tasks.

- [Reinventing your storage reality with IBM XIV Storage System e-Learning course](#)

Are you ready to reinvent the way your that business stores its data? In the next 5 minutes, you will learn how the IBM XIV Storage System solution is the only storage system that was designed by listening to our clients, just like you. The team at IBM listened to the challenges of leading enterprises in the marketplace and built a solution to solve these challenges. This solution includes high performance, reliable access to data, reduced costs of storing and managing data, and reduced complexity of managing storage environments and infrastructure.

- [Remote support e-Learning course](#)

The IBM XIV Storage System remote support feature, and its associated call home function, access the XIV Remote Support Center through the IBM XIV Storage System. In this lesson, you will learn about the remote support and the call home features.



Work-embedded e-Learning

Reinventing your storage reality with IBM XIV Storage System e-Learning course

Are you ready to reinvent the way your that business stores its data? In the next 5 minutes, you will learn how the IBM® XIV® Storage System solution is the only storage system that was designed by listening to our clients, just like you. The team at IBM listened to the challenges of leading enterprises in the marketplace and built a solution to solve these challenges. This solution includes high performance, reliable access to data, reduced costs of storing and managing data, and reduced complexity of managing storage environments and infrastructure.

The IBM XIV Storage System is designed to solve real challenges. The benefits of the design can make a tangible difference to ensuring the success of your business. The XIV system provides better data protection and availability than traditional products, true up-front and operational cost savings, a storage approach that serves your business processes and goals, and a new flexibility that can greatly increase your productivity.

You heard about what is possible with IBM XIV, but is it really that good? Do not take our word for it. Listen to a few real-world clients.

When the air conditioning failed in its primary data center, German electronics retailer Euronics was able to continue operations with zero downtime and no loss of data. In fact, no XIV customer has ever lost data.

Using IBM XIV storage, European trading center Luxembourg Stock Exchange reduced their total cost of ownership by 70%. They also reduced capital costs by a full 50% and electricity, footprint, and manpower costs by one-third.

International food producer Welch Foods implemented an XIV solution in an Oracle environment and, as a result, their transaction times were reduced by 60%. An innovative US healthcare facility, OCH Regional Medical Center, implemented an XIV solution to support their Tivoli® management software and VMware environment. They saw a 50% increase in performance and 70% - 80% increase in their speed of provisioning. BlueWater Power of Canada found that the self-managing features of an XIV solution provided them with consistently strong SAP performance, freeing their IT team to focus on supporting business development.

Thousands of other businesses, in every industry, have had similar experiences. They are finding that the IBM XIV system provides an exceptional infrastructure for meeting their organizational and business goals and is exceeding their expectations.

How does XIV storage do it? First, it provides unprecedented reliability and performance. Its revolutionary redundancy scheme includes a distributed data algorithm that stripes 1 MB blocks of data across two different disks in separate modules. When a disk fails, all disks in the system participate in the rebuild. The outcome? It takes less time for the XIV system to restore full redundancy. For example, if a disk on a fully used 1 TB drive system fails, XIV restores full redundancy in an average of 30 minutes or less, and without a significant performance penalty. Traditional RAID systems might take up to a day to restore data reliability, with performance compromised during rebuild. The fast XIV rebuild time means that your data spends less time that is exposed to risk. In this way, the XIV system actually protects your data more than traditional enterprise storage systems can. Also, every IBM XIV system provides business continuity through synchronous and asynchronous mirroring, and at no extra cost. So you can meet your disaster recovery needs flexibly and efficiently and know that your business will stay up and running thanks to non-disruptive access to all data at the secondary site. The IBM XIV Storage System solution delivers a lower total cost of ownership over competitors in several ways.

First, it greatly reduces the time and space that backups require. Traditional backup processes copy full volumes of data, including unused space, onto more capacity, and often with a performance penalty. The IBM XIV system uses its revolutionary built-in snapshot architecture to copy only the blocks of data that are written, enabling backup that has no performance hit, and is space- and time-efficient, and therefore less costly.

Second, the XIV system provides efficient capacity utilization and use of high-density disk drives to create a smaller footprint, and incur lower associated power and cooling costs than traditional storage products. In fact, the IBM XIV Storage System offers power and cooling savings of at least 63% over comparable traditional enterprise storage systems.

Third, managing XIV is simple, taking less time and opening the possibility of storage management to the broader IT team. And, the high reliability of the system means you can avoid the downtime costs that traditional storage typically incurs.

The IBM XIV Storage System is easy – a breakthrough in enterprise storage management. First, it is efficient. Its performance is high and consistent, so there are no hotspots to manage like in traditional storage systems.

Next, the intuitive management interface makes day-to-day storage administration simple. With the XIV system, it takes substantially less time to complete tasks, from provisioning storage, to seeing real-time capacity and performance metrics, and creating storage pools, not to mention setting up and replicating. And, the self-tuning and self-healing capabilities of the system greatly reduce the storage management activities common in traditional storage systems.

XIV systems also make tasks easier and more streamlined, with ready-to-go certified and supported integration solutions for your leading applications, such as email, databases, analytics, and virtualization, and with flexible options for connectivity and interoperability.

Finally, XIV native virtualized storage is the perfect fit with most virtualized server environments, serving as a key building block to your virtualization, and private cloud programs.

The IBM XIV Storage System has a proven track record of integrating with the industry-leading applications that your business relies on. These applications include virtualization technologies like VMware, Citrix XenServer, and Microsoft Hyper-V (TM); databases like Oracle, Microsoft SQL Server, and IBM DB2®; enterprise resource planning software like SAP and Oracle; email with Microsoft Exchange Server and IBM Lotus Notes®; and storage management and data protection solutions from IBM Tivoli, Symantec, and others.

Businesses are often challenged to do more with less. The solution is storage efficiency, which can be exceptionally achieved by using the IBM XIV Storage System.

Why look into XIV? The XIV system is the fastest growing storage solution in the industry today, with thousands of companies worldwide discovering how its unique architecture and benefits work to meet their growth challenges. And when you buy XIV, you are establishing a partnership with IBM, the worldwide IT powerhouse, offering end-to-end storage solutions for every need, and global support and services that are unmatched in the industry. And, the IBM XIV system is the optimal, cost-effective building block for your private cloud.

Today, you learned about the IBM XIV Storage System and its exceptional advantages, including reliability, reduced costs, and simplicity.

To learn more, contact your IBM representative, IBM Business Partner, or visit ibm.com®.



Work-embedded e-Learning

Remote support e-Learning course

The IBM® XIV® Storage System remote support feature, and its associated call home function, access the XIV Remote Support Center through the IBM XIV Storage System. In this lesson, you will learn about the remote support and the call home features.

The IBM XIV remote-support client is a software component inside the XIV system. It relies on a single outgoing TCP connection and cannot receive inbound communications. The client starts, terminates, and restarts all connections and is controlled using the command-line interface called the XCLI. Call home is an additional feature of remote support that must be set up at the initial IBM XIV Storage System installation by an IBM service support representative. With remote support, your XIV system can be repaired either remotely with the help of the XIV Remote Support Center or on-site by an IBM service support representative. When you contact an IBM remote-support specialist about a potential problem, that specialist can connect to your XIV system to analyze the problem and then decide whether it can be repaired remotely or on-site. If the XIV system is set up to work with the Support Center, you can have your IBM service support representative set alerts that are triggered when specific problems arise. The alerts automatically contact the XIV Remote Support Center whose specialists initiate appropriate action. Additionally, you can request that the same call-home alerts be sent to you.

Setting up your IBM XIV Storage System to connect directly to the XIV Remote Support Center takes a little planning. In this lesson, we review each step in detail.

The XIV Remote Support Center supports one of three methods of communication. The preferred method is a secure, high-speed, dedicated connection from the XIV system using a virtual private network or management port. Or you can choose a secure high speed broadband connection, or a modem and analog telephone line. Broadband Internet access is a dedicated connection providing high data-rate transmission to the Internet. Dial-up modems are limited to 56 kilobits per second and require the dedicated use of a telephone line. The first two dedicated connections require a Secure Shell to guarantee security. If the XIV system does not have direct access to the Internet (for example, due to a firewall), you can use the XIV Remote Support Proxy to facilitate the connection to IBM. You will learn more about the Remote Support Proxy later in this lesson.

You can initiate the connection to the XIV Remote Support Center through the IBM XIV Storage Management GUI or IBM XIV command-line interface (XCLI). In this lesson, you will learn how to create a connection using the IBM XIV Storage Management GUI. The GUI includes configuration steps to set up remote support and call home from your XIV system. Your IBM service representative performs these steps during your initial installation. After using your XIV system, you might decide to change some settings. Here are the steps needed to configure the remote support system.

From the GUI, click **Tools**, **Settings**, **Support**. The Support Configuration window is displayed. Your IBM service representative uses the Support Configuration screen to configure your ports. The XIV system supports both VPN and management ports. If you are not using broadband, determine whether you want to connect through a management or a VPN port. Obtain the appropriate IP addresses. You now have all the information your IBM service representative needs to complete your installation. The installer works through the Remote Support GUI, completing all five sections.

If you do not have a dedicated connection from your XIV system, a dedicated broadband connection or a telephone connection, you can use the XIV Remote Support Proxy to facilitate that connection. You can find the Remote Support Proxy information, including download, installation and configuration instructions, on the [IBM Support Portal website](#) on the Downloads window. For more information, consult the [IBM Support Portal website](#).

After your XIV system is installed, you are ready to use the remote support and call home features. It how this way: You initiate an Internet Secure Shell connection to the XIV Remote Support Center. The XIV Remote Support Center identifies your XIV system and marks it as connected. Your XIV system sends events to the database for analysis. If you allowed IBM to access your XIV system, the IBM service support representative logs on to the XIV Remote Support Center, takes over the call and connects to the XIV system being serviced. When service is complete, the IBM service support representative terminates the remote support session, and the XIV system disconnects from the XIV Remote Support Center. You can also terminate calls, or they can be stopped by an automatic timeout.

With call home, your XIV system can be set up to automatically contact the XIV Remote Support Center when triggered by an alert. Call home can make outgoing calls but cannot accept any incoming communication. It is designed to alert the IBM Support Center that there is a potential problem. The Support Center cannot contact your system. Call home is a component of the XIV Remote Support Center. The IBM XIV call home function connects your system directly to the IBM Trouble Ticketing System and automatically sends both failure and pre-failure notifications to the IBM Support Center. These notifications are sent through email from your SMTP server. After receiving the notifications at the Support Center, IBM service personnel can analyze the issues promptly and take appropriate action, including sending an IBM service representative to the site with a replacement part, if necessary. With access to the XIV system, IBM service personnel can

perform service tasks, such as viewing error logs and problem logs or initiating trace and memory dump retrievals. The call home configuration always includes SMTP, so you must have an SMTP email system available that the XIV system can use to send call home outgoing emails. Call home is never configured to provide incoming emails to the XIV system.

In this eLearning lesson, you learned about the XIV remote support feature. You also learned that you can choose from several configuration options such as the type of connection: dedicated or telephone, and the network ports. If those connection options are not available to you, you can download a Remote Support Proxy. Finally, you learned about the power of the IBM XIV Storage System call home function, where you can set up alerts that can be viewed by the XIV Remote Support Center. These alerts can trigger automatic service responses.

For more information, see [Introducing the IBM Remote Support Proxy utility](#).

Planning

This planning information defines preinstallation requirements for the IBM® XIV® Storage System Gen3.

Good planning is essential for the successful setup and use of your XIV system. It ensures that you have everything you need and that you meet all the prerequisites for the XIV system. It minimizes errors and helps installation go faster. Use the information provided to plan and prepare for placement of the XIV system, power and environmental needs, and unique configurations that are based on how you plan to use the XIV system.

Completing the Technical Delivery Assessment (TDA)

The [Technical Delivery Assessment \(TDA\) pre-installation checklist and worksheets](#) document provides several worksheets that the customer must complete for each XIV system. Use the worksheets to record information that is needed by the IBM installation planning representative (IPR) and IBM service representative to install and configure the XIV system.

Attention: Customers must prepare their environments to handle the XIV system based on this planning information with assistance from an IBM installation planning representative (IPR) or an IBM service representative. The final installation site within the computer room must be prepared *before* the equipment is delivered. If the site cannot be prepared before the delivery time, customers must make arrangements to have the professional movers return to finish the transportation later. Only professional movers can transport the equipment. The 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 devise alternative solutions.

IBM XIV Storage System Gen3 models

The following list briefly describes all current IBM XIV Storage System Gen 3 models:

- **281x-114:** The first Gen3 model.
- **281x-214:** The second Gen3 model, including hardware improvements over model 281x-114 and support for use of IBM Real-time Compression (RtC) as of XIV 11.6, available via licensing option.
- **281x-314:** The latest Gen3 model, with twice the RAM and CPU as 281x-214, optimized for the next level of IBM RtC, IBM Turbo Compression, included as a standard and pre-enabled feature.

Warranty periods

The following table lists the machine types and models (MTMs) that are available for the IBM XIV Gen3 Storage System, along with the warranty period for each.

Table 1. Available hardware and warranty periods for each

Machine type and model	Warranty
2810-114	One-year warranty
2812-114	Three-year warranty
2810-214	One-year warranty
2812-214	Three-year warranty
2810-314	One-year warranty
2812-314	Three-year warranty

- [System specifications](#)
This information compares general properties, performance, physical features, and host connectivity for each XIV Gen3 model.
- [Planning for compression](#)
Before implementing IBM Real-time Compression on your system, assess the current types of data and volumes that are used on your system.
- [Determining physical configuration and features](#)
Use these general guideline for determining and ordering the feature codes that you need to customize the XIV system.
- [Planning for delivery and relocation requirements](#)
Ensure that your environment meets the standard delivery clearance and weight requirements for the IBM XIV Storage System.
- [Planning for installation site requirements](#)
You must ensure that the location where you plan to install the IBM system meets all requirements.
- [Planning for network and host connectivity requirements](#)
This information describes the XIV system and host-SAN connections and requirements.
- [Planning for management-tool requirements](#)
The IBM XIV Storage System provides a set of management tools, including the IBM XIV Storage Management GUI, IBM Hyper-Scale Manager, IBM XIV online monitoring tool (XIVTop), and IBM XIV command-line interface (XCLI). You can use these management tools to configure and monitor your XIV systems. This information describes the hardware and software requirements and installation procedures for these management tools.

- [Planning for remote support and call home](#)

This information describes what you must know to plan for remote support and call home connections.

- [Planning for the rear-door heat exchanger](#)

If the optional rear-door heat exchanger (feature code 0082) was ordered, the customer must prepare the installation site before an IBM service representative can install the rear-door heat exchanger on any of your XIV systems.

System specifications

This information compares general properties, performance, physical features, and host connectivity for each XIV® Gen3 model.

General properties

Specification	Model 281x-114	Model 281x-214	Model 281x-314
Capacity per disk drive (7,200 RPM)	1 TB limited capacity, 2 TB, or 3 TB	1 TB limited capacity, 2 TB, 3 TB, 4 TB, or 6 TB	4 TB or 6 TB
Number of disk drives (minimum / maximum)	72 / 180	72 / 180	108 / 180
Capacity on-demand module configurations	Yes Note: Not available for 1 TB limited capacity systems.	Yes Note: Not available for 1 TB limited capacity systems.	Yes
Warranty	1 and 3 year limited warranty, onsite service, same day 24x7	1 and 3 year limited warranty, onsite service, same day 24x7	1 and 3 year limited warranty, onsite service, same day 24x7

Performance features

Specification	Model 281x-114	Model 281x-214	Model 281x-314
Maximum number of processors	15 Intel Xeon processor E5620	15 Intel Xeon processor E5645	30 Intel Xeon processors E5645
Maximum number of processor cores	60 physical (120 logical cores as a result of hyper-threading)	90 physical (180 logical cores as a result of hyper-threading)	180 physical (360 logical cores as a result of hyper-threading)
Maximum memory	360 GB (24 GB of memory per module)	For 1 TB, 2 TB, and 3 TB systems, 360 GB (24 GB of memory per module). For 4 TB, 6 TB systems, 720 GB (48 GB of memory per module).	1440 GB (96 GB of memory per module).
Maximum cache-to-disk bandwidth	480 Gbps	480 Gbps Note: For 4 TB or 6 TB options, 720 Gbps (48 GB per module)	720 Gbps
SSD caching (optional)	Up to 6.0 TB	Up to 12.0 TB	Up to 12.0 TB

Physical features

Specification	Model 281x-114	Model 281x-214	Model 281x-314
Dimensions (width x depth x height)	66 cm x 120 cm x 202 cm (79.5 in. x 26.0 in. x 47.2 in.)	66 cm x 120 cm x 202 cm (79.5 in. x 26.0 in. x 47.2 in.)	66 cm x 120 cm x 202 cm (79.5 in. x 26.0 in. x 47.2 in.)
Front clearance	120 cm (47.2 in.)	120 cm (47.2 in.)	120 cm (47.2 in.)
Rear clearance	100 cm (39.4 in.)	100 cm (39.4 in.)	100 cm (39.4 in.)
Maximum weight	1041.5 kg (2296.11 lbs)	1041.5 kg (2296.11 lbs)	1050 kg (2315.25 lbs)
Redundant power feed	Yes	Yes	Yes
Input voltage	180 - 264 V ac at 60 A or 30 A (±10%)	180 - 264 V ac at 60 A or 30 A (±10%)	180 - 264 V ac at 60 A or 30 A (±10%)
Power usage	See Power consumption	See Power consumption	See Power consumption

Temperature

Specification	Model 281x-114	Model 281x-214	Model 281x-314
Temperature	10 - 35° C (50 - 95° F)	10 - 35° C (50 - 95° F)	10 - 35° C (50 - 95° F)
Maximum altitude	2134 m (7000 ft.)	2134 m (7000 ft.)	2134 m (7000 ft.)
Humidity	25 - 80% non-condensing	25 - 80% non-condensing	25 - 80% non-condensing

Connectivity

Specification	Model 281x-114	Model 281x-214	Model 281x-314
Maximum number of Fibre Channel ports (capable of auto-negotiation)	24 (8 Gbps ports)	24 (8 Gbps ports)	24 (8 Gbps ports)
Fibre Channel rates	8 Gbps	8 Gbps	8 Gbps
Maximum number of iSCSI over Gigabit Ethernet ports	22 (1 Gbps Ethernet ports)	22 (1 Gbps Ethernet ports)	22 (1 Gbps Ethernet ports)
iSCSI rates	1 Gbps	1 Gbps or 10 Gbps	12 (10 Gbps Ethernet ports)

Planning for compression

Before implementing IBM Real-time Compression on your system, assess the current types of data and volumes that are used on your system.

Note: If you are using IBM XIV Storage System software version 11.6.x or later, and IBM XIV Management Tools version 4.6 or later, you can see storage saving estimates (even if compression is not licensed or is just disabled). The decision to compress volumes can be based on the expected storage savings of the compressed data and the expected effect on performance.

- [Understanding compression rates, ratios, and savings](#)
Consider a use case where the original capacity required to hold data was 100 TB, but 20 TB after compression (100 TB = 20 TB + 80 TB saved).
- [Prerequisites and limitations](#)
The following prerequisites and limitations are for XIV Storage System models 114 or 214, with XIV Storage software version 11.6 or later, and models 214 or 314 with XIV Storage software version 11.6.1 or later.
- [Estimating compression savings](#)
Compressible data can be identified and expected compression ratios can be estimated even before using compression.

Understanding compression rates, ratios, and savings

Consider a use case where the original capacity required to hold data was 100 TB, but 20 TB after compression (100 TB = 20 TB + 80 TB saved).

The following values help to clarify these terms:

Compression rate = 80%

Compression savings rate = 80%

Compression savings = 80 TB

Compression ratio = original size (100 TB) *divided by* the size on disk after compression (20 TB) = 5:1

When you consider savings, it is easiest to use the compression rate.

The compression ratio helps in understanding how much effective data you can store on your system. So, when you have a 5:1 compression ratio, you will be able to store 500 TB of data on 100 TB of physical capacity.

Prerequisites and limitations

The following prerequisites and limitations are for XIV Storage System models 114 or 214, with XIV Storage software version 11.6 or later, and models 214 or 314 with XIV Storage software version 11.6.1 or later.

Tip: Consider using version 11.6.1 with model 214 to benefit from improvements like compressed volume size.

- Compressed volumes must be created in thin-provisioned pools.
- To convert a mirrored volume from uncompressed to compressed (or vice versa), the mirroring relationship for that volume must first be removed and then recreated after the conversion.
- Snapshots cannot be converted from compressed to uncompressed and vice versa. Snapshots that already existed before a volume was converted from non-compressed to compressed are not converted and are not available with the converted volume.
- Space requirements:
 - Prior to enabling compression, the system must have a minimum of 17 GB of free hard space available. Enabling compression reserves 17 GB from the available system hard capacity. It is reserved for internal system use only.
 - Before the compression process, there must be enough space for both compressed and uncompressed versions of the volume.
 - For models 114 and 214 with version 11.6.0: Volume size must be at least 103 GB before compression.
 - For models 214 and 314 with version 11.6.1: Volume size must be at least 51 GB before compression.

The following is a partial list of limitations.

- Up to 1024 volumes and snapshots can be compressed.
- The following limits apply to compression capacity:
 - System must have a minimum of 17 GB of free hard space to enable IBM Real-time Compression.
 - Thin pools require a minimum of 17 GB of free hard space available to convert or transform volumes from uncompressed to compressed.
 - Thin pools require a minimum free soft space that is at least as large as the volume size that is being converted from uncompressed to compressed.
 - When you decompress a compressed volume, you must have both free hard space at least the size of the uncompressed volume and free soft space. It is a good practice to have free soft space at least the size of the uncompressed volume.
 - The Storage Admin can modify the system soft capacity. Use the XCLI command, `system_soft_capacity_set soft_size=SizeGB`. For example:

```
system_soft_capacity_set soft_size=250000
```

The soft capacity size of the system can be set to the following:

- Up to 3 times the size of the hard capacity of the system (with version 11.6.0)
- Up to 2 PB (with version 11.6.1)

And as low as the maximum size between the currently allocated soft capacity and the system's hard capacity (whichever is greater). The current hard, soft, and allocated soft capacity can be retrieved by using the **system_capacity_list** command.

Tip: Over-provisioning with Real-time Compression is safe, since compression ratios are predictable and stable.

- To compress an uncompressed volume in a thick pool, it must be moved to a thin-provisioned pool with compression enabled.
- Only one conversion process can be active at any time.
- Adding a module, rebuilding a disk, or upgrading the system suspends and then resumes the conversion process.

For the most current information about the limitations, refer to the [Limitations](#) section of the IBM XIV Storage System Gen3 Release Notes, versions 11.6.0 and 11.6.1.

Estimating compression savings

Compressible data can be identified and expected compression ratios can be estimated even before using compression.

On an XIV system supporting compression, the compression ratio for all uncompressed volumes in the system is continuously estimated, even before enabling compression. The decision to use compression can be based on the expected storage savings of the compressed data and the expected effect on performance (throughput and latency) of the compression processing overhead.

Information on compression usage can also be monitored using the XIV GUI to determine the potential savings to your storage capacity when uncompressed volumes are compressed. You can view the total percentage and total size of capacity savings when compression is used on the system. Compression savings across individual domains, pools and volumes can also be monitored. These compression values can be used to determine which volumes have achieved the highest compression savings. See the *IBM® XIV® Storage System Management Tools User Guide* for more information on monitoring and using compression.

Note: Keep in mind that the expected storage savings can vary from 5% higher or lower than the actual compression ratio. A negative estimated compression ratio can be due to metadata that consumes storage space on a volume, even when a 0% estimated compression ratio is received from data that cannot be compressed.

- **Effective capacity**
Effective capacity is the amount of storage that is allocated to applications.
- **General compression saving guidelines**
The best candidates for data compression are data types that are not already compressed. Compressed data types are used by many workloads and applications, such as databases, character/ASCII based data, email systems, server virtualization, CAD/CAM, software development systems, and vector data.
- **Estimating compression savings using XIV GUI**
The XIV software provides a built-in comprestimator function from the XIV Management Tools GUI on an XIV system supporting Real-time Compression.
- **Estimating compression savings using IBM Comprestimator utility**
Comprestimator is a stand-alone tool that can be used to estimate compression savings for data that is either not on XIV storage, or on an XIV Gen2 or Gen3 storage system with system software earlier than 11.6.x.

Effective capacity

Effective capacity is the amount of storage that is allocated to applications.

Using thin-provisioned storage architectures, the effective capacity is virtually larger than the array usable capacity. This is made possible by over-committing capacity, or by compressing the served data. Compression is the preferred method to apply thin-provisioning to usable capacity, since the compression ratio is highly predictable.

Hard capacity denotes usable, non-compressed capacity, whereas *soft capacity* denotes the nominal capacity that is assigned to volumes, and reported to any hosts mapped to those volumes. Thin provisioning denotes committing more soft capacity than hard capacity. Soft capacity is assigned at a pool level. In the case of compression, thin-provisioning is obvious: a compressed volume will forever use less hard capacity than soft capacity.

In XIV terms, the effective capacity is allocated out of the system soft capacity, and is the sum of the sizes of all the allocated volumes.

In an XIV System Gen3 Turbo Compression model 314, the maximal, non-compressed hard capacity supported in a single XIV frame is 485 TB (15 modules, 6 TB drives). However, an XIV frame can effectively accommodate up to 2 PB of real written data - when the data is compressed. With very high compression rates, filling a system up to 2 PB of soft capacity may not require a lot of usable capacity.

The maximum *soft capacity* that can be allocated to volumes in XIV is 2 PB. Considering the compression ratio for typical data profiles on XIV systems, the effective soft capacity leveraged within a 15-module frame will range from 1 PB to 2 PB. To maximize the utilization of XIV hard and soft capacity with compressed data, and avoid over-sizing the system, it is important to assess the expected compression ratio for the stored data. The Comprestimator command-line host-based utility can be used for that purpose. For more information on Comprestimator, see [Estimating compression savings using IBM Comprestimator utility](#).

The maximum effective capacity is reached when all the soft capacity has been allocated to volumes (that is, 2 PB).

For more information on how to fine-tune soft and hard capacities, refer to "Additional space utilization guidance" in the [Real-time Compression with IBM XIV Storage System Model 314 \(REDP-5306\)](#) Redbook.

General compression saving guidelines

The best candidates for data compression are data types that are not already compressed. Compressed data types are used by many workloads and applications, such as databases, character/ASCII based data, email systems, server virtualization, CAD/CAM, software development systems, and vector data.

The following examples represent workloads and data that are already compressed and are, therefore, not good candidates for compression:

- Compressed audio, video, and image file formats:
 - File types such as JPEG, PNG, MP3, medical imaging (DICOM), and MPEG2.
- Compressed user productivity file formats:
 - Microsoft Office 2007 and newer formats (.pptx, .docx, .xlsx, etc.), PDF files, Microsoft Windows executable files (.exe), etc.
- Compressed file formats:
 - File types such as .zip, .gzip, .rar, .cab, and .tgz.

IBM Real-time Compression is best suited for data that has an estimated compression savings of 25% or higher. There are various configuration items that affect the performance of compression on the system. Different data types have different compression ratios, and it is important to determine the compressible data currently on your system.

The IBM Comprestimator, a host-based utility, can be used to estimate expected compression rates. Compressing selectively, based on saving estimates, optimizes both capacity use and performance. For more information on Comprestimator, see [Estimating compression savings using IBM Comprestimator utility](#).

The following table shows the compression ratio for common applications and data types:

Table 1. Compression ratios for different data types

Data Types/Applications	Compression Ratios
Productivity	Up to 75%
Databases	Up to 80%
CAD/CAM	Up to 70%
Virtualization	Up to 75%

Note: The required capacity reserve is equal to the size of the volume (not the used capacity, but the volume size).

Estimating compression savings using XIV GUI


The XIV software provides a built-in comprestimator function from the XIV Management Tools GUI on an XIV system supporting Real-time Compression.

From the XIV GUI, compressible data can be identified and expected compression ratios can be estimated even before using compression. Compression does not even have to be enabled to view compression saving estimates. Continuous saving estimates are visible at all times for uncompressed volumes. And the compression ratio for all uncompressed volumes in the system is continuously estimated in a cyclical manner. That is, the potential savings estimations appear continuously and are updated every few hours.

The decision to use compression can be based on the expected storage savings of the compressed data and the expected effect on performance (throughput and latency) of the compression processing overhead.

[Figure 1](#) displays the compression savings (in both percentage and GB values) of compressed volumes and uncompressed volumes with *estimates* of potential savings, should the uncompressed volumes be compressed. These potential compression savings are constantly being updated.

Figure 1. Compression savings in the Volumes by Pools view



Compression Saving and Compression Saving (%) appear on the following views:

- Storage Pools
- Volumes by Pools
- Volumes and Snapshots
- Consistency Groups
- Domains
- Systems list

Estimating compression savings using IBM Comprestimator utility

Comprestimator is a stand-alone tool that can be used to estimate compression savings for data that is either not on XIV storage, or on an XIV Gen2 or Gen3 storage system with system software earlier than 11.6.x.

Comprestimator is a command-line host-based utility that can be used to estimate the expected compression rate for block-devices. The utility uses advanced mathematical and statistical algorithms to perform sampling and analysis efficiently. The utility also displays its accuracy level by showing the compression accuracy range of the results that are achieved based on the formulas it uses, deviating plus or minus 5 percent based on the formulas that are used by the RACE implementation.

The utility runs on a host that has access to the devices to be analyzed. It runs only read operations, so it has no effect on the data that is stored on the device. The following links provide useful information about installing Comprestimator on a host and using it to analyze devices on that host: [Comprestimator Utility](#) and [Comprestimator Utility Version 1.5.2.2](#).

For more information on Comprestimator, refer to the *IBM Real-time Compression on the IBM XIV Storage System (REDP-5215)* Redpaper at <http://www.redbooks.ibm.com/redpieces/abstracts/redp5215.html?Open&pdfbookmark>.

Determining physical configuration and features

Use these general guideline for determining and ordering the feature codes that you need to customize the XIV® system.

Procedure

To determine what features to order, answer the following questions:

1. How many XIV systems do you need?
2. Which XIV model configuration best fits your needs?
3. How many data and interface modules do you need?
4. What size disk drive do you need?
5. Does the installation site meet the physical site requirements for the XIV systems and features that you plan to order?

For example:

- Can you space the racks to allow for sufficient floor strength?
 - Is there adequate cooling capacity to handle the new equipment?
 - Is sufficient power available?
 - Do you need 30A or 60A power cables?
 - Is water cooling required?
- [Rack components](#)
This information lists the components that are supplied with each XIV system.
 - [Rack configuration features](#)
You can order an XIV system with a total of six or nine to 15 modules. Seven-module and eight-module configurations are not supported.
 - [Data and interface module features](#)
This information describes the data and interface module options for the IBM XIV Storage System.
 - [Solid-state drive caching features](#)
This information describes the solid-state drive (SSD) caching for the IBM XIV Storage System.
 - [Radio frequency identification device feature](#)
IBM offers an optional radio frequency identification device (feature code 0203) for the XIV system.
 - [Ruggedized rack feature](#)
The ruggedized rack (feature code 0080) is an optional kit for stabilizing the IBM XIV Storage System rack so that the rack complies with IBM earthquake-resistance standards. The ruggedized rack secures the rack directly to the concrete floor and stiffens the rack from flexing. The kit can be installed on a raised floor or non-raised floor.
 - [Rear-door heat exchanger feature](#)
The rear-door heat exchanger (feature code 0082) is an optional water-cooled device that is mounted on the rear of an IBM XIV Storage System. It cools the air that is heated and exhausted by devices inside the rack.
 - [Weight-and-height-reduced shipping feature](#)
This information describes the Weight-and-height-reduced shipping features for Weight-and-height-reduced shipping.

Rack components

This information lists the components that are supplied with each XIV® system.

Models 281x-114, 281x-214, and 281x-314

The following components are supplied with each XIV system:

- IBM® T42 rack
- For models 281x-114 and 281x-214:
 - Three to nine data modules, each with 12 disk drive modules (DDMs) and one optional solid-state drive (SSD)
 - Three or six interface modules, each with 12 DDMs and one optional SSD
- For model 281x-314:
 - Three to nine data modules, each with 12 DDMs and one SSD (standard/required)
 - Six interface modules, each with 12 DDMs and one SSD (standard/required)
 - A minimum of nine total modules (three data and six interface)
- Three uninterruptible power supplies
- Two InfiniBand switches
- One maintenance module

- One automatic transfer switch
- One modem
Note: The modem is no longer supported as part of the rack configuration for Indonesia and Brazil. For more information, contact your IBM planning representative.
- Six host-connections patch panels with Fibre Channel and Ethernet (iSCSI) ports
- Utility patch panel with maintenance, technician, management, remote support, and modem ports
- Complete set of internal cabling

Rack configuration features

You can order an XIV® system with a total of six or nine to 15 modules. Seven-module and eight-module configurations are not supported.

Note: For model 281x-314, a minimum total of nine modules are required, three data and six interface. For more information, see [Rack components](#). [Table 1](#) shows the various system configuration options with their feature codes.

Note: The six module feature does not apply to model 281x-314.

Table 1. Feature codes for rack configurations

Feature description	Feature code
6 module initial capacity	0806
9 module initial capacity	0809
10 module initial capacity	0810
11 module initial capacity	0811
12 module initial capacity	0812
13 module initial capacity	0813
14 module initial capacity	0814
15 module initial capacity	0815

The IBM® XIV Storage System requires a corresponding software license for each entitled data or interface module.

Partially populated configurations have the same power, floor planning, hardware, software, and cabling configurations as fully populated systems. In other words, partially populated racks are identical to fully populated rack, except that they have fewer modules and therefore less storage capacity.

Note: XIV systems with less than 13 modules also have fewer usable Fibre Channel and iSCSI ports. For more information, see [Table 1](#) in [Connectivity](#).

Fully populated rack (15 modules)

A fully populated rack contains nine data modules and six interface modules, for a total of 15 modules.

Each module contains 12 1 TB limited capacity, 2 TB, 3 TB, 4 TB, or 6 TB disk drive modules (DDMs), for a total of 180 DDMs in a fully populated rack (15 modules x 12 DDMs).


Note: Only the 4 TB and 6 TB DDM options are available for model 281x-314. [Table 2](#) shows the various configuration limits for a fully populated rack's data storage.

Table 2. Fully populated rack limits for data storage

Hard drive size	Total usable capacity
1 TB	81 TB
2 TB	161 TB
3 TB	243 TB
4 TB	325 TB
6 TB	485 TB

The following figure shows an example of the front and rear views of a fully populated XIV system.

Figure 1. Example of a fully-populated XIV system model 281x-114, model 281x-214, or model 281x-314



Note: The modem is no longer supported as part of the rack configuration for Indonesia and Brazil. For more information, contact your IBM planning representative.

Minimally populated rack

A minimally populated rack for models 281x-114 and 281x-214 contains three data modules and three interface modules, for a total of six modules. A minimally populated rack for model 281x-314 contains three data modules and six interface modules, for a total of nine modules.

Each module contains twelve 1 TB limited capacity, 2 TB, 3 TB, or 4 TB, 3 TB, 4 TB, or 6 TB disk drive modules (DDMs), for a total of 72 DDMs (6 modules x 12 DDMs) for models 281x-114 and 281x-214 and 108 DDMs (9 modules x 12 DDMs) for model 281x-314.

Note: Only the 4 TB and 6 TB DDM options are available for model 281x-314.


[Table 3](#) shows the various minimal configuration limits for a rack's data storage.

Table 3. Minimally populated rack limits for data storage

Hard drive size	Total usable capacity for models 281x-114 and 281x-214	Total usable capacity for model 281x-314
1 TB	28 TB	N/A
2 TB	55 TB	N/A
3 TB	84 TB	N/A
4 TB	112 TB	177 TB
6 TB	169 TB	267 TB


The following figure shows an example of the front and rear views of a minimally populated XIV system for models 281x-114 and 281x-214.

Figure 2. Example of a minimally populated XIV system



Note: The modem is no longer supported as part of the rack configuration for Indonesia and Brazil. For more information, contact your IBM planning representative. The next figure shows an example of the front and rear views of a minimally populated XIV system for model 281x-314.

Figure 3. Example of a minimally populated XIV system for model 281x-314



Note: The modem is no longer supported as part of the rack configuration for Indonesia and Brazil. For more information, contact your IBM planning representative.

- **Capacity**
Each data module and interface module contains twelve 1 TB limited capacity, 2 TB, 3 TB, 4 TB, or 6 TB disk drive modules (DDMs).
- **Processors and memory**
The information lists the number of processors and memory that are available in fully populated and partially populated racks.
- **Connectivity**
This information lists the number of Fibre Channel and Ethernet (iSCSI) ports that are available in fully and partially populated racks.

Capacity

Each data module and interface module contains twelve 1 TB limited capacity, 2 TB, 3 TB, 4 TB, or 6 TB disk drive modules (DDMs).

You can only order an XIV® system with DDMs of the same capacity, either 1 TB hard drive size, 2 TB, 3 TB, 4 TB, or 6 TB. You cannot order an XIV system with a mix of DDM capacities.

Note: Only the 4 TB and 6 TB DDM options are available for model 281x-314.

For data and interface modules with 1 TB limited-capacity, the modules contain 2 TB DDMs, but the usable capacity is limited to 1 TB by using the *short-stroking method*. By short stroking the DDM, only the outer sectors of the disk platters are used to store data. Head repositioning delays are minimized, which results in shorter access times and increased I/O performance.

The following tables list the total capacity and usable capacity available in fully populated and partially populated racks for each supported DDM capacity. The usable capacity is the amount of space that can be used for data storage. The remaining space is used for system activities, data redundancy, and spare capacity. For usable capacity, the decimal is rounded down.

Table 1. Capacity for models 281x-114 and 281x-214 XIV systems

Number of modules in the rack	Number of interface modules	Number of data modules	Number of disk drives	Total usable capacity in TB (1 TB limited capacity / 2 TB / 3 TB / 4 TB / 6 TB)
6	3	3	72	28 / 55 / 84 / 112 / 169

Number of modules in the rack	Number of interface modules	Number of data modules	Number of disk drives	Total usable capacity in TB (1 TB limited capacity / 2 TB / 3 TB / 4 TB / 6 TB)
9	6	3	108	44 / 88 / 132 / 177 / 267
10	6	4	120	51 / 102 / 154 / 207 / 311
11	6	5	132	56 / 111 / 168 / 225 / 338
12	6	6	144	63 / 125 / 190 / 254 / 382
13	6	7	156	67 / 134 / 203 / 272 / 409
14	6	8	168	75 / 149 / 225 / 301 / 453
15	6	9	180	80 / 161 / 243 / 325 / 485

Table 2. Capacity for model 281x-314 XIV systems

Number of modules in the rack	Number of interface modules	Number of data modules	Number of disk drives	Total usable capacity in TB (4 TB / 6 TB)
9	6	3	108	177 / 267
10	6	4	120	207 / 311
11	6	5	132	225 / 338
12	6	6	144	254 / 382
13	6	7	156	272 / 409
14	6	8	168	301 / 453
15	6	9	180	325 / 485

Note: DDMs in the XIV system do not need to be repaired immediately upon failure because the XIV system is fully redundant, even with the failures. Replacement for failed DDMs is deferred until the third DDM fails, or until another component requires immediate onsite maintenance. For administrators, if a DDM fails but replacement is deferred, the IBM® XIV Storage Management GUI displays the DDM as pale red, and the Status as Failed (Deferred). When there are three or more DDMs in failed state, the IBM XIV Storage Management GUI displays the failed DDMs as bright red, and the Status as Failed. For IBM service representative (technicians), the IBM XIV Storage Management GUI displays the failed DDM as bright red, and the Status as Failed regardless of whether the replacement is deferred.

Processors and memory

The information lists the number of processors and memory that are available in fully populated and partially populated racks.

Table 1. Processors and memory for XIV model 281x-314

Number of modules in the rack	Number of processors	Memory cache in GB (4 TB / 6 TB)	SSD cache in TB (800 GB SSD)
9	18	864	7.2
10	20	960	8.0
11	22	1056	8.8
12	24	1152	9.6
13	26	1248	10.4
14	28	1344	11.2
15	30	1440	12

Connectivity

This information lists the number of Fibre Channel and Ethernet (iSCSI) ports that are available in fully and partially populated racks.

The maximum number of iSCSI host ports that can be logged-in concurrently is 700 per interface module. The maximum number of Fibre Channel (FC) host ports that can be logged-in concurrently is 350 per target port and 700 per interface module. However, the limit of 700 per interface module includes FC and iSCSI host ports logged-in concurrently, and mirror connectivity. The per system limit is 1400 concurrent port log-ins for a six module system and 2000 concurrent port log-ins for nine or more module systems (including FC, iSCSI and mirror connectivity). This applies to XIV Gen3 microcode version 11.1.1 or later.

The following table lists the number of Fibre Channel and iSCSI ports and identifies the interface modules from which Fibre Channel and iSCSI ports are available for each module configuration. For XIV® systems with 1 Gb Ethernet ports, interface module 4 has four Fibre Channel ports and two Ethernet (iSCSI) ports. Interface modules 5 - 9 have four Fibre Channel ports and four Ethernet (iSCSI) ports each. For XIV systems with 10 Gb Ethernet ports, all interface modules have four Fibre Channel ports and two Ethernet (iSCSI) ports.

Table 1. Fibre Channel and iSCSI ports on interface modules

Number of modules in the rack	Interface modules with Fibre Channel and iSCSI ports	Number of Fibre Channel ports	Number of iSCSI ports (1 Gb ³ / 10 Gb)
6	4, 5	8 ^{1, 2}	6 / 4
9	4, 5, 7, 8	16 ^{1, 2}	14 / 8
10	4, 5, 7, 8	16 ^{1, 2}	14 / 8
11	4, 5, 7, 8, 9	20 ¹	18 / 10
12	4, 5, 7, 8, 9	20 ¹	18 / 10
13	4, 5, 6, 7, 8, 9	24	22 / 12
14	4, 5, 6, 7, 8, 9	24	22 / 12
15	4, 5, 6, 7, 8, 9	24	22 / 12

Number of modules in the rack	Interface modules with Fibre Channel and iSCSI ports	Number of Fibre Channel ports	Number of iSCSI ports (1 Gb ³ / 10 Gb)
Notes:			
1. Fibre Channel and iSCSI ports on module 6 are not active in configurations of 12 or fewer modules. 2. Fibre Channel and iSCSI ports on module 9 are not active in configurations of 10 or fewer modules. 3. On interface module 4, only two of the four available Ethernet ports are used.			

Data and interface module features

This information describes the data and interface module options for the IBM® XIV® Storage System.

For data and interface modules with 1 TB limited-capacity, the modules contain 2 TB DDMs, but the usable capacity is limited to 1 TB by using the *short-stroking method*. By short stroking the DDM, only the outer sectors of the disk platters are used to store data. Head repositioning delays are minimized, which results in shorter access times and increased I/O performance.

Capacity on Demand (CoD) modules are modules that are physically installed in the rack and fully functional but that are not 'activated' for use until more capacity is needed. Additional capacity is activated by ordering a miscellaneous equipment specification (MES) with the appropriate CoD activation feature. In reality, installing the CoD activation feature allows for the use of an additional, incremental portion of total system capacity that is proportional to the share of a CoD module in the system. For example in a 15-module system, if three of the modules are ordered as CoD (and not yet activated), this means that initially the customer is only licensed to use 12/15 of the system usable capacity. By activating the three CoD modules, the customer will then be licensed to fully utilize the system usable capacity (or 15/15 of the capacity).

CoD activation features do not require physical installation. If you receive notification to install this feature, only report the installation according to the field-reporting guide for your respective geographical area.

Table 1. Feature codes for data modules

Feature description	Feature code		
	Model 281x-114	Model 281x-214	Model 281x-314
1 TB limited-capacity data modules	1123	2016	
2 TB data modules	1126	2026	
2 TB CoD data module	1146	2028	
2 TB module CoD activation	1149	4020	
2 TB self-encrypting module CoD activation		4021	
3 TB data modules	1136	2036	
3 TB CoD data module	1148	2038	
3 TB module CoD activation	1150	4030	
3 TB self-encrypting module CoD activation		4031	
4 TB data modules		2042	5042
4 TB CoD data module		2044	5044
4 TB self-encrypting module CoD activation		4041	6041
6 TB data modules		2062	5062
6 TB CoD data module		2064	5064
6 TB self-encrypting module CoD activation		4061	6061

Table 2. Feature codes for interface modules

Feature description	Feature code		
	Model 281x-114	Model 281x-214	Model 281x-314
1 TB limited-capacity interface modules with 1 GbE cards	1122	2015	
1 TB limited capacity interface modules with 10 GbE cards		2115	
2 TB interface modules with 1 GbE cards	1125	2025	
2 TB interface modules with 10 GbE cards		2125	
2 TB CoD interface module with 1 GbE cards	1145	2027	
2 TB CoD interface module with 10 GbE cards		2127	
2 TB module CoD activation with 1 GbE cards	1149	2029	
2 TB module CoD activation with 10 GbE cards		4120	
2 TB self-encrypting module CoD activation with 10 GbE cards		4121	
3 TB interface modules with 1 GbE cards	1135	2035	
3 TB interface modules with 10 GbE cards		2135	
3 TB CoD interface module with 1 GbE cards	1147	2037	
3 TB CoD interface module with 10 GbE cards		2137	
3 TB module CoD activation with 1 GbE cards	1150	2039	
3 TB module CoD activation with 10 GbE cards		4130	
3 TB self-encrypting module CoD activation with 10 GbE cards		4131	
4 TB self-encrypting interface modules with 1 GbE cards		2041	5041
4 TB self-encrypting interface modules with 10 GbE cards		2141	5141
4 TB self-encrypting CoD interface module with 1 GbE cards		2043	5043

Feature description	Feature code		
	Model 281x-114	Model 281x-214	Model 281x-314
4 TB self-encrypting CoD interface module with 10 GbE cards		2143	5143
4 TB self-encrypting module CoD activation with 1 GbE cards		2049	6041
4 TB self-encrypting module CoD activation with 10 GbE cards		4141	6141
6 TB self-encrypting interface modules with 1 GbE cards		2061	5061
6 TB self-encrypting interface modules with 10 GbE cards		2161	5161
6 TB self-encrypting CoD interface module with 1 GbE cards		2063	5063
6 TB self-encrypting CoD interface module with 10 GbE cards		2163	5163
6 TB self-encrypting module CoD activation with 1GbE cards		4061	6061
6 TB self-encrypting module CoD activation with 10GbE cards		4161	6161

Solid-state drive caching features

This information describes the solid-state drive (SSD) caching for the IBM® XIV® Storage System.

SSDs, also known as *flash drives*, are storage devices that use nonvolatile solid-state memory to reduce the latency and time to access the stored data. IBM offers an SSD caching option, which uses SSDs for caching across all volumes. SSD caching improves application performance by dynamically and adaptively caching the most frequently accessed data on SSDs.

This feature includes a 400, 480, or 800 GB SSD in each data and interface module, for up to 12.0 TB of management-free SSD caching.

The following table lists the feature available for each XIV model.

Note: Model 281x-314 includes 800 GB SSDs as standard features in each data and interface module used in the system, so there are no SSD caching feature codes associated with model 281x-314.

Table 1. Solid-state-drive caching features

Feature description	Feature code	
	Model 281x-114	Model 281x-214
400 GB SSD caching	1155	2255
480 GB SSD caching		2255
800 GB SSD caching		2265
Note: 800 GB SSD caching is mandatory for 281x-214 when the either the 4 TB or 6 TB option is used.		

Radio frequency identification device feature

IBM® offers an optional radio frequency identification device (feature code 0203) for the XIV® system.

If you use frequency identification device (RFID) technology to track equipment in your data centers, you can order the RFID option (feature code 0203) to attach an RFID tag on an XIV system. This RFID is designed to meet the performance and numbering specification as outlined by the radio frequency identification specifications. For information about the specification, see the [Financial Services Technology Consortium website](#).

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 XIV system that you want to track. This option does not tag individual components.

Important: If the tag must be replaced for an XIV system, ensure that you update the asset-management database with the new RFID number for that XIV system.

Ruggedized rack feature

The ruggedized rack (feature code 0080) is an optional kit for stabilizing the IBM® XIV® Storage System rack so that the rack complies with IBM earthquake-resistance standards. The ruggedized rack secures the rack directly to the concrete floor and stiffens the rack from flexing. The kit can be installed on a raised floor or non-raised floor.

It is important for XIV systems to be adequately restrained during earthquakes to prevent human injury and limit potential damage to critical system components, such as hard disk drives. The optional ruggedized rack includes hardware that secures the rack to the floor and x-braces on the front and rear of the rack, which prevents the rack from twisting. The earthquake-resistance option brings the XIV system into compliance with earthquake-resistance objectives that are documented in *Earthquake Resistance for IBM Hardware Products* (IBM Corporate Bulletin C-B 1-9711-009 9202).

Installation of the required floor hardware and the earthquake resistance kit is disruptive. Before you install the earthquake resistance kit, the XIV system must be turned off and temporarily moved while the floor is prepared and the kit is installed.

Notes:

- When you order the optional ruggedized rack, order one for each XIV system rack.
- To remove a previously installed ruggedized rack, contact your IBM service representative to remove the kit.

The rack tie downs are intended for securing a frame weighing 1,134 kg (2500 lbs) per rack. These tie downs are designed to secure the rack on either a non-raised floor or a raised floor installation. Instructions, drawings, and data are provided on an "as is" basis, WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OR CONDITIONS OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. IBM has made every effort to conduct limited tests, but not all situations are tested, and the drawings and data are provided on an "AS IS" basis, with no warranty of any kind, express or implied. Rather, IBM has provided the information to help procure the parts that are needed and either install the design or request IBM to install the design as a service.

The following table lists the feature availability for each XIV model.

Table 1. Ruggedized rack features

Feature description	Feature code		
	Model 281x-114	Model 281x-214	Model 281x-314
Ruggedized rack	0080	0080	0080

See [Preparing for the ruggedized rack](#) for information about requirement for preparing the installation site before the ruggedized rack option can be installed.

Rear-door heat exchanger feature

The rear-door heat exchanger (feature code 0082) is an optional water-cooled device that is mounted on the rear of an IBM® XIV® Storage 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. 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



The following table lists the feature availability for each XIV model.

Table 1. Rear-door heat exchanger features

Feature description	Feature code		
	Model 281x-114	Model 281x-214	Model 281x-314
Rear-door heat exchanger	0082	0082	0082

See [Planning for the rear-door heat exchanger](#) for information about requirements for preparing the installation site before the rear-door heat exchanger feature can be installed.

Weight-and-height-reduced shipping feature

This information describes the Weight-and-height-reduced shipping features for Weight-and-height-reduced shipping.

IBM® offers Weight-and-height-reduced shipping (feature code 0200) for XIV® systems that are ordered with 12 to 15 modules. This option reduces the weight of the rack by removing data modules 12 - 15 and shipping them separately. Each module with all disk drive modules installed weighs approximately 28.6 kg (63 lb). This option reduces weight loading on floors and elevators that are used for moving the unit from the truck to the final position in the installation site. The IBM service representative installs the components that were shipped separately during system installation.

A fully configured XIV model 281x-114, model 281x-214, and model 281x-314, including packaging, weighs approximately 1132.5 kg (2491.5 lb). This option reduces the weight to approximately 1018 kg (2239.5 lb).

If your site does not meet the delivery clearances that are shown in [Planning for delivery and relocation requirements](#), the weight-and-height-reduced shipping option must be ordered to reduce the height of the rack by 38 cm (15 in.). After the rack is delivered, the IBM service representative removes the rack top cover so that the XIV system 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.

Note: This option greatly increases the system installation time. Onsite coordination is needed for both the IBM service representative and the professional movers. Attention: Do not tilt the XIV system rack more than 12.5 degrees.

Table 1. Shipping features

Feature description	Feature code		
	Model 281x-114	Model 281x-214	Model 281x-314
Weight-and-height-reduced shipping	0200	0200	0200

Planning for delivery and relocation requirements

Ensure that your environment meets the standard delivery clearance and weight requirements for the IBM® XIV® Storage 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.

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 XIV Storage System as close to its final destination as possible. You must ensure that the loading ramp and receiving area can accommodate the XIV shipment.
- [Shipment weights and dimensions](#)
To help you plan for the delivery of the IBM XIV Storage System, ensure that the loading dock and receiving area can support the weight and dimensions of the packaged XIV shipments.

Planning to receive delivery

The professional movers or riggers are responsible for delivering and unloading the IBM® XIV® Storage System as close to its final destination as possible. You must ensure that the loading ramp and receiving area can accommodate the XIV shipment.

Procedure


Use the following steps to ensure that the receiving area and loading ramp can safely accommodate the delivery of your XIV 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 XIV 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 must be reduced for delivery to locations where doorway heights are smaller than the delivery clearances, the weight-and-height-reduced shipping (feature code 0200) must be ordered. This option reduces the pallet to 178 cm (70.0 in.). For information about the weight and height, see [Weight-and-height-reduced shipping feature](#).

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 12.5 degrees, as shown in the following figure.

Figure 1. Maximum tilt for a packaged rack is 12.5 degrees



A ramp with a maximum angle of 12.5 degrees must not be higher than 6.6 cm (2.6 inches) vertically for each 30.5 cm (12 in.) of horizontal length.

Shipment weights and dimensions

To help you plan for the delivery of the IBM® XIV® Storage System, ensure that the loading dock and receiving area can support the weight and dimensions of the packaged XIV shipments.

For models 281x-114, 281x-214, and 281x-314, at least one shipping container is delivered for each XIV system that you order. The container is a wooden pallet that is covered by a corrugated fiberboard (cardboard). This container contains the XIV system rack with components installed, and other items such as, power cords, CDs and printed pubs, and other features or peripherals for your model. Additional containers are delivered for optional features such as the ruggedized rack, rear-door heat exchanger.

Delivery clearance requirements

The clearance measurements (height x width x 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	Model 281x-114, 281x-214, and 281x-314
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. The following delivery weights are for a packaged rack with 15 modules. For model 281x-114 and 281x-214 racks with less than 15 modules, subtract the 28.6 kg (63 lb) for each module less than 15 to get the approximate weight requirement for your rack. For model 281x-314 racks with less than 15 modules, subtract the 29.7 kg (65.49 lb) for each module less than 15 to get the approximate weight requirement for your rack.

- **Model 281x-114:** 1132.5 kg (2491.5 lb)
- **Model 281x-214:** 1132.5 kg (2491.5 lb)
- **Model 281x-314:** 1150 kg (2535.75 lb)

Planning for installation site requirements

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

Plan your installation site with assistance from an IBM installation planning representative (IPR) or an IBM service representative. 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. Only professional movers can transport the equipment.

An IBM service representative installs the XIV® 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 XIV system rack as close to the installation site as possible because of its weight. Professional movers or riggers are also required to relocate or dispose of the XIV system.

- [Planning for floor and space requirements](#)
You must ensure that the location of the XIV system meets space and floor requirements.
- [Planning for power requirements](#)
Ensure that your operating environment meets the ac-power and voltage requirements.
- [Planning for environmental requirements](#)
Ensure that the installation site meets all operating environment requirements for the XIV system.

Planning for floor and space requirements

You must ensure that the location of the XIV® system meets space and floor requirements.

Procedure

Complete the following steps to ensure that the planned installation location meets space and floor load requirements:

1. Decide whether the XIV system is to be installed on a raised floor.
2. Determine whether the floor meets the floor-load requirements for the XIV system.
3. Calculate the amount of space need for the rack footprint 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.
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.
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.

- **[Raised or non-raised floor considerations](#)**

You can install the IBM XIV Storage System 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 IBM XIV Storage System.

- **[Rack dimensions and service clearance requirements](#)**

The installation site must accommodate the rack dimensions and minimum service clearance for the XIV 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 host cables, network Ethernet cable, and customer analog phone line.

- **[Preparing for the ruggedized rack](#)**

If the optional ruggedized rack (feature code 0080) was ordered, the customer must purchase fastening hardware and prepare the location where the kit is to be installed before an IBM service representative can install the ruggedized rack on any of your XIV systems.

- **[Preparing for the rear-door heat exchanger](#)**

If the optional rear-door heat exchanger (feature code 0082) was ordered, the customer must prepare the installation site before an IBM service representative can install the rear-door heat exchanger on any of your XIV systems.

Raised or non-raised floor considerations

You can install the IBM® XIV® Storage System 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 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 plyon) 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 to ensure efficient cooling for the XIV system; however, overhead cabling at the rear of the rack is available when the XIV 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 IBM® XIV® Storage System.

Floor reinforcement must support the weight of the XIV system over a specific area, as shown in the following table. These measurements are slightly less than the footprint area of the system, due to the overhang of the door.

The weight requirements in the following table are for a rack with 15 modules. For model 281x-114 and 281x-214 racks with less than 15 modules, subtract the 28.6 kg (63 lb) for each module less than 15 to get the approximate weight requirement for your rack. For model 281x-314 racks with less than 15 modules, subtract the 29.7 kg (65.49 lb) for each module less than 15 to get the approximate weight requirement for your rack.

Table 1. Floor weight-support requirements

Model	Floor reinforcement area	Total weight
Model 281x-114	66 cm × 118 cm (26 in. × 46.5 in.)	1040.8 kg (2294.5 lbs)
Model 281x-214	66 cm × 118 cm (26 in. × 46.5 in.)	1040.8 kg (2294.5 lbs)
Model 281x-314	66 cm × 118 cm (26 in. × 46.5 in.)	1050 kg (2535.75 lbs)

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 feature](#).

To ensure that all requirements are met, obtain the service of a qualified structural engineer to prepare the floor.

Important: 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 XIV® system.

IBM® service representative must have enough space to open the front and rear covers to service the XIV system, including removing components and other assemblies from the XIV system.


Notes:

- 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.

Dimensions and minimum service clearance for the XIV system are shown in the following table and figure.

Table 1. Rack dimensions and clearance requirements

Dimension	Model 281x-114, Model 281x-214, and Model 281x-314
Height	202 cm (79.5 in.)
Depth	120 cm (47.2 in.)
Width	66 cm (26.0 in.)
Front clearance	120 cm (47.2 in.)
Rear clearance	100 cm (39.4 in.)
Side (door) clearance	10 cm (3.9 in.)



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 XIV® system based on raised floor with 60 cm × 60 cm (24 in. × 24 in.) tiles:

1. Based on your planned layout, ensure that the installation site can accommodate the locations of the cables exiting each XIV system.
2. Plan for the XIV system to be positioned on two tiles, with the rear of rack aligned on a floor-tile seam and with a two full rows of perforated tiles immediately in front of the rack (see [Figure 1](#)).

Figure 1. Raised floor requirements



3. Cut a 200 mm × 200 mm (8 in. × 8 in.) opening in the rear floor tile for under-floor cabling and electricity (see [Figure 1](#)).
Note: Place the opening under the XIV system, centered on the back edge of the tile along the rear of the rack.
Sizing the cutout correctly is important. 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 XIV system (see [Figure 1](#)). These tiles must have a minimum of 40% perforation.
5. If the ruggedized rack is ordered, see [Ruggedized rack feature](#) for additional floor preparation steps.
6. 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 host cables, network Ethernet cable, and customer analog phone line.

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.

Notes:

- 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 ruggedized rack is ordered, see [Ruggedized rack feature](#) for more floor preparation steps.

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 ruggedized rack

If the optional ruggedized rack (feature code 0080) was ordered, the customer must purchase fastening hardware and prepare the location where the kit is to be installed before an IBM service representative can install the ruggedized rack on any of your XIV systems.


Procedure

The customer must perform one of the following procedures before the ruggedized rack installation, depending upon whether the XIV system sets on a raised or a non-raised floor. For either type of installation, work with a consultant or structural engineer to ensure that the site preparations meet the requirements.

- **Raised floor with a depth of 228.6 - 558.8 mm (9 - 22 in.)**

1. Consult the manufacturer of the floor panels that are used at the installation site to ensure that the floor panels can sustain a concentration load of 454 kg (1000 lb) per rack. More panel supports (pedestals) might be required to restore structural integrity of the panels. For multiple rack installation, two casters might induce a concentration load as high as 907 kg (2000 lb) on a floor panel.
2. Obtain the service of a qualified consultant or structural engineer to determine the appropriate heavy-duty concrete or slab-floor eyebolts. Each eyebolt must withstand 2268 kg (5000 lb) pull force per eyebolt.
3. Purchase four 0.5x13-in. heavy-duty concrete or slab eyebolts, with an internal diameter of 1-3/16 in.
4. Install the eyebolt fasteners 527.8 mm (20.8 in.) side to side and 965 mm (38 in.) front to back.

Figure 1. Layout and dimensions for circular cutouts



Important: The eyebolts must be installed in the concrete floor beneath the raised floor.

The height from the floor surface to the center of the internal diameter of the eyebolts must be 25.4 - 63.5 mm (1 - 2.5 in.) (see [Figure 2](#)). Higher than 63.5 mm (2.5 in.) might cause excessive lateral deflection to the tie-down hardware.

Figure 2. Eyebolts installed in the subfloor



- **Non-raised (concrete) floor**

1. Obtain the service of a qualified consultant or structural engineer to determine the appropriate concrete anchors to attach the rack-mounting plates to the concrete floor. Each anchor must withstand 2268 kg (5000 lb) of pull force.
2. Purchase eight concrete anchors.

Preparing for the rear-door heat exchanger

If the optional rear-door heat exchanger (feature code 0082) was ordered, the customer must prepare the installation site before an IBM® service representative can install the rear-door heat exchanger on any of your XIV® systems.

Procedure

To complete the rear-door heat exchanger site preparation, see [Planning for the rear-door heat exchanger](#).

Planning for power requirements

Ensure that your operating environment meets the ac-power and voltage requirements.

The IBM® XIV® Storage System contains three uninterruptible power supplies and redundant main-power cables that maintain power to the XIV system in the event of an ac-power loss.

The three uninterruptible power supplies in the XIV system are recognized by the central uninterruptible power supply system in the main data center as standard modern storage-system power supplies. The XIV system does not impose any harmful characteristics to the central uninterruptible power supply in the main data center.

The XIV system has redundant main-power cables. For two-main-power-cable configuration, you must supply power from two independent sources of electricity. Consult an IBM service representative to discuss power source options for the four-main-power-cable configuration.

Note: Removing all ac power from the XIV system causes an emergency shutdown. All modified data is then saved to disks, and the XIV system turns off within 5 minutes.

Customer responsibilities

- You must supply enough branch circuits to prevent overloading from the equipment that you install.
- You must ensure that each electrical outlet is correctly wired and grounded to prevent an electrical shock.

IBM responsibilities

- The IBM service representative complete several checks, including voltage and grounding checks before the power to the XIV 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 outlets are required for the two power cords that are needed by each XIV system.

Important: To eliminate a single point of failure, the outlets must be independent. This means that each outlet must use a separate power source and each power source must have its own wall circuit breaker.

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 XIV systems. The XIV 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.

- [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 IBM XIV Storage System.
- [Main-power cables specifications](#)
Plan for power cable, connector, and receptacle requirements.
- [Emergency power-off considerations](#)
The IBM XIV Storage System does not have an emergency power-off (EPO) switch; however, you can connect the XIV system to a room EPO switch.

Power sources

Several ac power source configurations are available.

- Two 30 A, 200-240 V ac, three-phase receptacles, each connected to a different power source
- Two 60 A, 200-240 V ac, three-phase receptacles, each connected to a different power source
- Two 60 A, 200-240 V ac, single-phase receptacles, each connected to a different power source.

The IBM® XIV® Storage System is protected by an uninterruptible power supply for internal use. 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.

J1 is the *primary feed*. This allows for better data center power usage because the ATS returns to J1 if there is power to that connector.

If the three-phase, 30A ATS for non-EMEA (feature code 9894) is installed and both power feeds into the XIV system have voltages within the operational tolerance, the power load is shared across both power feeds. If a power feed fails, the ATS switches the full power load to the remaining feed.

To prepare the power cables, you must provide your IBM service representative with the following information before the installation:

- The distance from the rack to the power receptacles.
- Whether overhead or under-floor power-cabling is to be used.

Power consumption

This information describes the power consumption for partial and full rack configurations.

The following table lists the power consumption for each supported disk-drive module (DDM) capacity with solid-state drives (SSDs). If no SSDs are present, the power consumption is marginally less.

The power consumption for partial-rack configurations is the typical value and is provided only for reference purposes. To support concurrent capacity upgrades, the installation site must provide sufficient power capacity to support high performance on a full rack configuration of 15 modules.

The high-performance measurements in the following tables were taken in an environment with a room temperature of 25° C (77° F).

Table 1. Power consumption

Number of modules in the rack	Number of disk drives	Model 281x-114 with SSDs in kVA High performance (1 and 2 TB / 3 TB)	Model 281x-214 with SSDs in kVA High performance (2 TB / 3 TB / 4 TB / 6 TB)	Model 281x-314 with SSDs in kVA High performance (4 TB / 6 TB)
6	72	2.9 / 2.9	2.5 / 2.6 / 2.6 / 2.6	N/A
9	108	4.3 / 4.2	3.6 / 3.8 / 3.9 / 3.9	4.7 / 4.7
10	120	4.7 / 4.7	4.0 / 4.2 / 4.3 / 4.3	5.1 / 5.1
11	132	5.1 / 5.2	4.3 / 4.5 / 4.6 / 4.6	5.5 / 5.5
12	144	5.5 / 5.7	4.7 / 4.9 / 5.0 / 5.0	6.0 / 6.0
13	156	5.9 / 6.1	5.0 / 5.3 / 5.4 / 5.4	6.5 / 6.5
14	168	6.5 / 6.7	5.5 / 5.7 / 5.8 / 5.8	7.0 / 7.0
15	180	6.9 / 7.2	5.8 / 6.1 / 6.2 / 6.2	7.5 / 7.5

Input voltages and frequencies

This information lists the voltages or frequencies that are accepted by the IBM® XIV® Storage System.

[Table 1](#) provides the voltages and frequencies that are accepted by the XIV system. For more information, see [Leakage current](#).

Table 1. Input voltages and frequencies

Characteristics	Voltage or frequency
Nominal input voltages	200, 220, or 240 V ac
Maximum input current	30 A or 60 A
Minimum input voltage	180 V ac
Maximum input voltage	264 V ac
Maximum ground-leakage current	0.3 A
Steady-state input frequencies	50±3 Hz or 60±3 Hz
Power line disturbance (PLD) input frequencies	50±3 Hz or 60±3 Hz

Main-power cables specifications

Plan for power cable, connector, and receptacle requirements.

The XIV® systems are available in multiple main-power cable configurations. The automatic transfer switch (ATS) is chosen based on the required main-power cables.

Attention: The regional designations are general. The ATS assigned to the regional designation might not be accurate for your country. If the input voltage for the country uses a wye circuit, use the appropriate ATS and 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 ATS and main-power cables for United States, Canada, Latin America. For more information about electric currents for various countries, see the [International Trade Administration website](#).

Feature codes

The following table provides feature codes and descriptions for each available ATS. It also lists the feature codes for each of the main-power cables that are associated with that specific ATS.

Table 1. Automatic transfer switches


Automatic transfer switch feature code	Description	Feature codes for associated main-power cables
9890	ATS for worldwide use; single-phase, 60 A	9811 9813
9891	ATS for United States, Canada, Latin America, and Japan; three-phase delta, 60 A	9820
9892	ATS for Europe, Middle East, Asia/Pacific (except Japan), Australia, and New Zealand; three-phase wye, 30 A	9822
9894	ATS for United States, Canada, Latin America, and Japan; three-phase delta, 30 A	9824

The following table provides the feature codes and description details for each main-power cable.

Note: The main-power cables are provided with the connector part numbers listed in the following table. Receptacle part numbers in the following table are recommended. Although equivalent receptacles can be used, it is the responsibility of the customer to verify compatibility.

Table 2. Main-power cables

Main-power cable feature code	ATS feature code	Description	Connector	Wall receptacle	Inline receptacle

Main-power cable feature code	ATS feature code	Description	Connector	Wall receptacle	Inline receptacle
9811	9890	Main-power cable for United States, Canada, Latin America, and Japan; single-phase, 60 A, two wires + ground with plug 	Hubbell HBL360P6V04	Hubbell HBL360R6W	Hubbell HBL360C6W
9813	9890	Main-power cable for Europe, Middle East, and Asia/Pacific (except Japan); single-phase, 60 A, two wires + ground, hard wired	Hard wired or uses client-provided connector	Not applicable	Not applicable
9820	9891	Main-power cable for United States, Canada, Latin America, and Japan; three-phase delta, 60 A, three wires + ground with plug 	Hubbell HBL460P9V05	Hubbell HBL460R9W	Hubbell HBL460C9W
9822	9892	Main-power cable for Europe, Middle East, Asia/Pacific (except Japan), Australia, and New Zealand; three-phase wye, 30 A, three wires + neutral + ground, hard wired	Hard wired or uses client-provided connector	Not applicable	Not applicable
9824	9894	Main-power cable for United States, Canada, Latin America, and Japan; three-phase delta, 30A, three wires + ground with plug 	Hubbell HBL430P9V04	Hubbell HBL430R9W	Hubbell HBL430C9W

Main-power cable considerations

The following information provides details to consider regarding main-power cables for IBM® System Storage® systems:

- Main-power cables are 250 V 60 A or 250 V 30 A rated. Single-phase main-power cables have two poles 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 9811 and 9820 are IEC309-compliant.
- Bulk wire conductor sizes for main-power cable feature codes are:
 - 9811 and 9820: 6 AWG
 - 9813: 10 mm²
 - 9822: 6 mm²
 - 9824: 10 AWG
- 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.
- Some feature codes have specific wall circuit-breaker requirements:
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.
 - Single-phase installations (200 V - 240 V) with ATS feature code 9890 and two main-power cables require wall circuit breakers that have a rating of 60 A to 63 A.
 - Three-phase delta installations (200 V - 240 V) with ATS feature code 9891 require wall circuit breakers that have a rating of 50 A to 60 A.
 - Three-phase wye installations (220 V - 240 V) with ATS feature code 9892 require wall circuit breakers that have a rating of 30 A to 32 A.
 - Three-phase delta installations (200 V - 240 V) with ATS feature code 9894 require wall circuit breakers that have a rating of 30 A.

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 XIV systems. The XIV 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. For more information, see [Leakage current](#).
- Main-power cables with the following feature codes must meet specific design certifications:
 - 9891 uses bulk wire that meets UL and CSA safety agency certifications, but not the EU harmonization certification.
 - 9892 uses bulk wire that meets the EU harmonization certification, but not UL and CSA safety agency certifications.
 - 9894 uses bulk wire that meets UL and CSA safety agency certifications, but not the EU harmonization certification.

Emergency power-off considerations

The IBM® XIV® Storage System does not have an emergency power-off (EPO) switch; however, you can connect the XIV system to a room EPO switch.

Note: National or local building, electrical, fire prevention, safety, and other laws or regulations might address or control the manner in which information technology equipment is installed within certain facilities and environments. The application of those laws or regulations might depend on considerations of factors beyond the nature or design of the equipment to be installed. It is a customer responsibility to interpret and identify any laws or regulations applicable to the installation of information technology in its environment and to inform IBM, IBM Business Partners, or their designated installers of any actions not identified in this planning guide that are necessary to install information technology equipment in the customer facilities in accordance with such applicable laws or regulations. A room EPO switch can be powered by the XIV system uninterruptible power supplies for use with non-powered switch circuit breakers. The EPO circuit in the uninterruptible power supplies is considered a Class 2 circuit (UL, CSA standards) and a safety extra-low voltage (SELV) circuit (IEC standard). Contact your IBM service representative for more information about connecting to a room EPO switch.

Attention: Powering off the XIV system using a room EPO switch results in data loss and possible loss of configuration. An IBM service representative is required to recover an XIV system that was turned off using a room EPO switch. If the XIV system loses ac power but is not powered off using an EPO circuit, data and configuration are preserved.

Planning for environmental requirements

Ensure that the installation site meets all operating environment requirements for the XIV® 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 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 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 XIV 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 XIV system.
 - [Operating vibration requirements](#)
The vibration levels that are designed for the XIV system comply with class V1L requirements included in the product classes for vibration.

Operating 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 XIV® system must meet the following requirements.

Operating (powered on)

- Temperature: 10 - 35°C (50 - 95°F)
- Relative humidity: 20 - 80%, non-condensing
- Maximum wet bulb temperature: 25°C (77°F)
- Maximum altitude: 2134 m (7000 ft)

Non-operating (powered off)

- Temperature: 5 - 45°C (41 - 113°F)
- Relative humidity: 20 - 80%, non-condensing
- Maximum wet bulb temperature: 27°C (80.6°F)

Shipping

- Temperature: -40 - 60°C (-40 - 140°F)
- Relative humidity: 5 - 95%, non-condensing
- Maximum wet bulb temperature: 29°C (84.2°F)

Important: An XIV system that runs continuously must be within the specified operating environment. An XIV 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, and with the air temperature at the front of the rack less than 27°C (80.6°F).

- [System environmental acclimation](#)

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 XIV® system.

Procedure

To optimize the cooling around your XIV system, complete the following steps:

1. Install the XIV 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 raised-floor installation and cabling](#).

3. (Model 281x-114, model 281x-214, and model 281x-314 only) Consider using the rear-door heat exchanger (feature code 0082) to further 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 feature](#).

4. Ensure that the installation site meets the cooling (thermal dissipation) requirements that are listed in the following table.
The following table lists thermal-dissipation requirements for an XIV system with 15 modules. To support future capacity upgrades, the installation site must provide cooling arrangements to support a full 15-module rack configuration.

Table 1. Thermal dissipation

Model type	Thermal dissipation (kBTU/hour)		
	1 TB limited capacity, or 2 TB	3 TB	4 TB or 6 TB
281x-114	22.7	24.0	
281x-214	19.1	20.1	20.5
281x-314	N/A	N/A	26

5. Ensure that the installation site meets the airflow requirements that are listed in the following table.
The airflow for partial-rack configurations is the typical value and is provided only for reference purposes. To support future capacity upgrades, the site must provide more airflow space to support a full-rack configuration.

The airflow values assume a maximum inlet air temperature of 25°C (77°F).

Table 2. Airflow requirements

Number of modules	Cubic feet per minute (CFM)	
	1 TB limited capacity, or 2 TB	3 TB, 4 TB, or 6 TB
6	300	310
9	440	450
10	480	500
11	520	550
12	560	600
13	600	650
14	640	700
15	690	750

- [Temperature threshold and events](#)

Temperature threshold and events

Attention: The XIV Storage System handles overheating by informing the administrator through warning events and initiating an automatic thermal shutdown as a last resort.

Note:

- 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. The following tables depict the temperature thresholds and events received. For systems running microcode version 11.6 or later, see [Table 1](#). For systems running microcode version 11.1 – 11.5, see [Table 2](#).
- 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. Temperature thresholds and events for systems running microcode version 11.6 or later

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	Check whether there is a problem that affects the data center temperature or the system temperature.
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.
45°C (113°F)	SYSTEM_TEMPERATURE_IS_CRITICALLY_HIGH_SHUTTING_DOWN	If automatic shutdown is enabled: Automatic system shutdown is in progress.

Table 2. Temperature thresholds and events for systems running microcode versions 11.0–11.5

Temperature threshold	Event	Description
29°C (80°F)	SYSTEM_TEMPERATURE_IS_OK_NOW	No action required.
33°C (82.4°F)	SYSTEM_TEMPERATURE_HIGH	System should be closely monitored and, if possible, take action to cool down.
35°C (95°F)	SYSTEM_TEMPERATURE_IS_TOO_HIGH	Actions to cool down the system must be taken.
45°C (113°F)	SYSTEM_TEMPERATURE_IS_CRITICALLY_HIGH	Manually shutdown the system by using the shutdown command.

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 IBM® XIV® Storage System hardware. Risks that are posed by the presence of excessive particulate levels or concentrations of harmful gases include damage that might cause the XIV 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 XIV system, IBM might require implementation of appropriate remedial measures to mitigate such environmental contamination before providing repair or replacement of the XIV 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 71.04-1985¹, which states that the reactivity rate of copper coupons must be less than 300 Angstroms per month (Å/month, $\approx 0.0039 \mu\text{g}/\text{cm}^2\text{-hour weight gain}$)². In addition, the reactivity rate of silver coupons must be less than 300 Å/month ($\approx 0.0035 \mu\text{g}/\text{cm}^2\text{-hour weight gain}$)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% RH⁴. Data centers must be free of zinc whiskers⁵.

¹ ANSI/ISA-71.04.1985. *Environmental conditions for process measurement and control systems: Airborne contaminants*. Instrument Society of America, Research Triangle Park, NC, 1985.

² The derivation of the equivalence between the rate of copper corrosion product thickness growth in Å/month and the rate of weight gain assumes that Cu₂S and Cu₂O grow in equal proportions.

³ The derivation of the equivalence between the rate of silver corrosion product thickness growth in Å/month and the rate of weight gain assumes that Ag₂S is the only corrosion product.

⁴ The 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 XIV® 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 an XIV system while either the front or rear door is open or when the front and rear doors are not installed.

The acoustic levels are shown in [Table 1](#). These measurements are based on a typical XIV system configuration with 13 modules. All measurements are in conformance with ISO 7779 and declared in conformance with ISO 9296.

Notes:

1. LWAd is the statistical upper-limit A-weighted sound-power level (rounded to the nearest 0.1 B).
2. LpAm is the mean A-weighted emission sound-pressure level that is measured at the 1-meter bystander positions (rounded to the nearest dB).
3. 10 dB (decibel) = 1 B (bel)

Table 1. Acoustic declaration

	Declared A-weighted sound power level, LWAd (B)	Declared A-weighted sound pressure level, LpAm (dB)
Operating	7.5	58
Idle	7.5	58

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, compliance with such government regulations also depends upon a variety of additional factors, including the duration of employee exposure and whether employees wear hearing protection. IBM® recommends that you consult with qualified experts in this field to determine whether you are in compliance with the applicable regulations.

Operating vibration requirements

The vibration levels that are designed for the XIV® system comply with class V1L requirements included in the product classes for vibration.

The XIV system is designed to operate under the vibration V1L levels that are described in the following table. More information includes random vibration PSD profile breakpoints and operational shock levels.

Table 1. Vibration levels

Class	g _{rms}	g Peak Sine
V1L	0.10	0.06 @ 50 & 60 Hz

Notes:

- All values in this table are in g²/Hz.
- g is the peak g level of an approximate half-sine pulse.

Table 2. Random vibration PSD profile breakpoints

Class	5 Hz	17 Hz	500 Hz
V1L	2.0 x 10 ⁻⁷	2.2 x 10 ⁻⁵	2.2 x 10 ⁻⁵

Note: All values in this table are in g²/Hz.

Table 3. Operational shock levels

Class	Axis	g ¹	pw ²
1	Vertical	3.5	3.0

Notes:

- g is the peak g level of an approximate half-sine pulse.
- pw is the pulse width in milliseconds.

Planning for network and host connectivity requirements

This information describes the XIV® system and host-SAN connections and requirements.

- [Network connections](#)
This information describes the patch panel and utility panels, which are used for network connectivity for IBM XIV Storage System.
- [Management ports requirements](#)
The management ports are connected to systems that are used for managing the XIV system through the IBM XIV Storage Management GUI and IBM XIV command-line interface (XCLI). These ports can also be used for sending e-mail notifications and SNMP traps about event alerts.
- [Network configurations](#)
Host systems can connect to an XIV system over a Fibre Channel network 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 connections

This information describes the patch panel and utility panels, which are used for network connectivity for IBM® XIV® Storage System.

Important: 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.

Host-connections patch panel

Host systems are connected to the XIV system through a SAN network by using Fibre Channel or Ethernet (iSCSI) ports in the host-connections patch panel. There is a host-connections patch panel for each interface module.

The following figures illustrate the host-connections patch panels for XIV systems with 1 Gb Ethernet ports or 10 Gb Ethernet ports.

Figure 1. Host-connections patch panel with Fibre Channel and 10 Gb Ethernet iSCSI ports





Figure 2. Host-connections patch panel with Fibre Channel and 1 Gb Ethernet iSCSI ports



The maximum number of iSCSI host ports that can be logged-in concurrently is 700 per interface module. The maximum number of Fibre Channel (FC) host ports that can be logged-in concurrently is 350 per target port and 700 per interface module. However, the limit of 700 per interface module includes FC and iSCSI host ports logged-in concurrently, and mirror connectivity. The per system limit is 1400 concurrent port log-ins for a six module system and 2000 concurrent port log-ins for nine or more module systems (including FC, iSCSI and mirror connectivity). This applies to XIV Gen3 microcode version 11.1.1 or later.

The following table lists the number of Fibre Channel and iSCSI ports and identifies the interface modules from which Fibre Channel and iSCSI ports are available for each module configuration. For XIV systems with 1 Gb Ethernet ports, interface module 4 has four Fibre Channel ports and two Ethernet (iSCSI) ports. Interface modules 5 - 9 have four Fibre Channel ports and four Ethernet (iSCSI) ports each. For XIV systems with 10 Gb Ethernet ports, all interface modules have four Fibre Channel ports and two Ethernet (iSCSI) ports.

Table 1. Fibre Channel and iSCSI ports on interface modules

Number of modules in the rack	Interface modules with Fibre Channel and iSCSI ports	Number of Fibre Channel ports	Number of iSCSI ports (1 Gb ³ / 10 Gb)
6	4, 5	8 ^{1, 2}	6 / 4
9	4, 5, 7, 8	16 ^{1, 2}	14 / 8
10	4, 5, 7, 8	16 ^{1, 2}	14 / 8
11	4, 5, 7, 8, 9	20 ¹	18 / 10
12	4, 5, 7, 8, 9	20 ¹	18 / 10
13	4, 5, 6, 7, 8, 9	24	22 / 12
14	4, 5, 6, 7, 8, 9	24	22 / 12
15	4, 5, 6, 7, 8, 9	24	22 / 12


Notes:

1. Fibre Channel and iSCSI ports on module 6 are not active in configurations of 12 or fewer modules.
2. Fibre Channel and iSCSI ports on module 9 are not active in configurations of 10 or fewer modules.
3. On interface module 4, only two of the four available Ethernet ports are used.

Utility patch panel

The following table describes the ports that are available in the patch panel and the component to which each patch-panel port connects. A diagram of the utility patch panel is also shown in the "Patch panel" column.

Table 2. Utility patch panel connections

Patch panel	Patch panel port	Component	Description
	Maint (maintenance module) ports	Maintenance module	Use these two ports to connect directly to the maintenance module.
	Management ports	Modules 1, 2, 3	Use these three redundant ports to connect to systems that are used for managing the XIV system using the IBM XIV Storage Management GUI and IBM XIV command-line interface (XCLI). These ports can also be used for sending email notifications and SNMP traps about event alerts, and communication with key servers for managing encryption keys. Note: Use all three connections for redundancy. If the customer uses only one path and the module for that connection goes down, the ability to manage the XIV system is lost. For information about IP configuration and protocol requirements, see Management ports requirements .
	VPN (remote support) ports	Modules 1, 3	Use these two virtual private network (VPN) ports are used to connect to the XIV Remote Support Center. Note: Use both VPN connections for redundancy to avoid losing remote support if one of the modules fails.
	Tech (technician) ports	(Model 281x-11x only) Modules 2, 4 (Models 281x-21x and 281x-314 only) Modules 2, 3	Use these two ports to connect the XIV system directly to a notebook computer on which the GUI and XCLI are installed. The IBM service representative uses these ports to initial configure and service the XIV 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 XIV system. To establish a connection between the notebook computer and the XIV system, you must manually set the Ethernet adapter on the notebook computer to use a fixed address and subnet mask. Do not use the Dynamic Host Configuration Protocol (DHCP) server that is implemented over this port. Important: - The technician port must be used only by an IBM service representative. It is <i>not</i> intended for customer use. - Use the technician ports only for a direct-connection to a single notebook-computer. Do <i>not</i> connect the technician ports to the network. A connection to the network might cause incorrect IP address assignments to the host systems. - The technician port for module 4 is active on XIV system code V11.0.1 or later. The port is not active on XIV system code V11.0.0. (Model 218x-11x only) The technician port that is labeled "3" is logically routed to module 3 through module 4.
	Modem port	Modem	Use this port to connect to a modem for remote support. Note: The modem is no longer supported as part of the rack configuration for Indonesia and Brazil. For more information, contact your IBM planning representative.

Management ports requirements

The management ports are connected to systems that are used for managing the XIV® system through the IBM® XIV Storage Management GUI and IBM XIV command-line interface (XCLI). These ports can also be used for sending e-mail notifications and SNMP traps about event alerts.

To ensure redundancy in case of interface-module failure, the IBM XIV Storage System is managed through three IP addresses over Ethernet interfaces. Each IP address is handled by a different interface module. The IBM XIV Storage Management GUI and XCLI must be configured with all three IP addresses that are defined for the XIV system. Management functions can be performed through any of the IP addresses. These addresses can be accessed simultaneously by multiple clients. Users must configure the IBM XIV Storage Management GUI or XCLI for the set of IP addresses that are defined for the specific XIV system.

To be resilient to a network failure, ensure that the network equipment that provides management communication is protected by an uninterruptible power supply.

IP configuration

For each of the three management ports, the customer or technical assistant must provide the following information to the IBM service representative before the installation is started:

- IP address of the port
- Net mask
- Default IP gateway
- Maximum transmission unit (MTU)

Note: All management IP interfaces must be connected to the same subnet and use the same network mask, gateway, and MTU.

In addition, provide the following system-level IP information:

- IP address of the primary and secondary DNS servers
- IP address or DNS names of the SNMP manager, if required
- IP addresses or DNS names of the SMTP servers

Protocols

The XIV system is managed through dedicated management ports that are running TCP/IP over an Ethernet connection. Management is carried out through the following protocols:

- The XIV system communicates with key servers using the KMIP protocol over TCP port 5696 (if the encryption feature is enabled).
- Proprietary protocols are used to manage the XIV system from the IBM XIV Storage Management GUI and IBM XIV command-line interface (XCLI). This management communication is over TCP port 7778. The XIV GUI and XCLI act as the client and initiate the connection. The XIV system acts as the server. All traffic over this port is encrypted through the Secure Sockets Layer (SSL) protocol.
- The XIV system responds to SNMP requests when sending replies to SNMP managers.
- The XIV system initiates SNMP packets when sending traps to SNMP managers.
- The XIV system initiates SMTP traffic when sending emails for either event notifications or for SMS gateways. SMTP connections use port 25.
- The IBM XIV Storage Management GUI and XCLI communication is Secure Sockets Layer (SSL) encrypted, which provides security for passwords and various information.
- The management ports must be on a different subnet than the VPN ports that are used for remote access.

Network configurations

Host systems can connect to an XIV® system over a Fibre Channel network or over an Ethernet network by using the Internet Small Computer System Interface (iSCSI) protocol.

Restriction: Direct attachment between hosts and the XIV system is not supported.

Host systems must not have a mix of Fibre Channel and iSCSI connections to the XIV system. If your situation requires you to temporarily mix Fibre Channel and iSCSI host access to migrate data, you must plan carefully to avoid data loss.

Important:

- Host traffic can be directed to any of the interface modules. The storage administrator must ensure that host connections avoid single points of failure. The server administrator must ensure that the host workload is adequately balanced across the connections and interface modules. This balancing can be done by installing the relevant host-attachment kit. Review the balancing periodically and when traffic patterns change.
- Always employ multipathing for data connections. If multipathing is not used, a host-connection error is displayed. If there is a problem on one path in the connection and an interface module fails, a host-connection error is also displayed.
Not all host systems support multipath configuration between the protocols.
- Host systems must have the appropriate host-attachment kit installed. Host attachment kits are available for AIX®, HP-UX, Linux®, Solaris, and Microsoft Windows operating systems. For more information, see [IBM® XIV Host Attachment Kit](#). For supported interoperability configurations, see the [System Storage® Interoperation Center website](#). Also see the [Host Attachment and Interoperability RedBook](#).
- [Fibre Channel network configurations](#)
Host systems can connect to an XIV system over a Fibre Channel network. In a production environment, always connect Fibre Channel hosts to a minimum of two separate SAN switches in independent fabrics to provide redundancy.
- [Ethernet \(iSCSI\) network configurations](#)
Host systems can connect to an XIV system over an Ethernet network using the Internet Small Computer System Interface (iSCSI) protocol.

Fibre Channel network configurations

Host systems can connect to an XIV® system over a Fibre Channel network. In a production environment, always connect Fibre Channel hosts to a minimum of two separate SAN switches in independent fabrics to provide redundancy.

Several network configurations that use Fibre Channel are technically possible, and each configuration varies in terms of cost, flexibility, performance, and reliability.

Three redundant configurations are described: 12 paths to each volume, 6 paths to each volume, and minimal cabling. Each of these configurations has no single point of failure:

- If an interface module fails, each host remains connected to all other interface modules.
- If a Fibre Channel switch fails, each host remains connected to at least three interface modules.
- If a host bus adapter (HBA) on a host fails, the host remains connected to at least three interface modules.
- If a host cable fails, the host remains connected to at least three interface modules.

The following example configurations are for an XIV system over a Fibre Channel with six interface modules. For an XIV system over a Fibre Channel with three interface modules, fewer paths are used, but the principle of balancing paths across all interface module still applies.

The minimum supported configuration is paths to two different interface modules. One path to a single interface module is not supported.

Host system can have as many HBAs as needed to support the operating system, application, and overall performance requirements.

In an XIV system, each interface module has four Fibre Channel ports.

For supported interoperability configurations, see the [System Storage® Interoperation Center website](#).

Redundant configuration with six paths to each volume


[Figure 1](#) illustrates a redundant configuration in which all six interface modules are accessed by using six paths per LUN on the host. This configuration is recommended for most general-purpose configurations. Host reliability during path error recovery on certain operating systems is complicated by increasing the

number of paths per LUN. For host systems with two HBAs, the six paths per LUN method is recommended.

In this configuration:

- Each host is equipped with dual HBAs. Each HBA (or HBA port) is connected to one of two Fibre Channel switches.
- Each Fibre Channel switch has a connection to a separate Fibre Channel port on each of the six interface modules.
- One host uses the first three paths per fabric, and the next host uses the three other paths per fabric.
- If a fabric fails, all interface modules are still used.
- Each volume has six paths.

Figure 1. Fully redundant Fibre Channel configuration with six paths to each volume




Redundant configuration with 12 paths to each volume

Figure 2 illustrates a redundant configuration in which all six interface modules are accessed by using 12 paths per LUN on the host. This configuration is ideal for benchmarking and higher performance host systems with the highest I/O requirements.

In this configuration:

- Each host is equipped with dual HBAs. Each HBA (or HBA port) is connected to one of two Fibre Channel switches.
- Each Fibre Channel switch has a connection to a separate Fibre Channel port on each of the six interface modules.
- Each volume can be accessed through 12 paths. Twelve paths is the maximum number of paths to consider. There is no benefit to having more than 12 paths. Using more than 12 paths might cause issues with host processor usages and server reliability if a path failure occurs (such as a host HBA failure).

Figure 2. Fully redundant Fibre Channel configuration with 12 paths to each volume




Redundant configuration with minimum cabling and paths to all interface modules

Figure 3 illustrates a redundant configuration with minimum cabling and paths from each host LUN to all six interface modules. In this configuration:

- Each host is equipped with dual HBAs. Each HBA (or HBA port) is connected to one of two Fibre Channel switches.

- Each Fibre Channel switch has a connection to three separate interface modules.
- Each volume has six paths.

Figure 3. Fully redundant Fibre Channel configuration with minimum cabling and paths to six interface modules



Ethernet (iSCSI) network configurations

Host systems can connect to an XIV® system over an Ethernet network using the Internet Small Computer System Interface (iSCSI) protocol.

Several network configurations using Ethernet are technically possible, and each configuration varies in terms of cost, flexibility, performance, and reliability.

In the IBM® XIV Storage 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)

Notes:

1. If two or more iSCSI connections exist in your iSCSI network, iSCSI ports for mirroring must be defined on switches with Flow Control Send On.
2. MTU configuration is required if the network supports an MTU that is greater than the default 1536 bytes. Specify the largest possible MTU.
For XIV system code earlier than 11.2.0.a, the maximum MTU is 4500 bytes.

For XIV system code 11.2.0.a and later, the maximum MTU is 9000 bytes.

3. Because the XIV 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 XIV system.
4. If present, Ethernet VLANs and IP routers must be configured to enable connectivity between the host systems and the XIV system.
5. If present, IP routers must be configured to enable access between the hosts and the XIV system.

Important: Link aggregation is not supported. Ports cannot be bonded.

For supported interoperability configurations, see the [System Storage® Interoperation Center website](#).

Redundant configuration

In this configuration, each interface module is connected through two ports to two different Gigabit Ethernet switches, and each host system is connected to the two Gigabit Ethernet switches. Full connectivity is maintained in the event of a component failure.

This configuration is the same as the redundant configuration options for Fibre Channel networks.

For the best performance, use a dedicated iSCSI network infrastructure.

Single-switch configuration

In this configuration, a single Gigabit Ethernet switch interconnects all host systems and interface modules.

Use non-redundant configurations only when the risks of a single point of failure are acceptable, which is typically the case for test and development environments.

Single-port host configuration

In this configuration, a single Gigabit Ethernet switch connects to each host system and to two interface modules.

Use non-redundant configurations only when the risks of a single point of failure are acceptable, which is typically the case for test and development environments.

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. Required cable types

Cables	Required cable type
Fibre Channel	50 μ m (micrometer) multimode Fibre Channel cables with LC/LC or LC/SC connectors Note: If you require 62.5- μ m fibers, contact an IBM® service representative for assistance.
Ethernet (iSCSI)	For model A14 or 114/214/314 with 1 Gb, straight copper gigabit CAT5e-rated Ethernet cables with RJ-45 connectors or For model 114/214/314 with 10 Gb, 50 μ m (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
Modem	CAT5 telephone cable with RJ-11 connectors Note: The modem is no longer supported as part of the rack configuration for Indonesia and Brazil. For more information, contact your IBM planning representative.

Planning for management-tool requirements

The IBM® XIV® Storage System provides a set of management tools, including the IBM XIV Storage Management GUI, IBM Hyper-Scale Manager, IBM XIV online monitoring tool (XIVTop), and IBM XIV command-line interface (XCLI). You can use these management tools to configure and monitor your XIV systems. This information describes the hardware and software requirements and installation procedures for these management tools.

- **Requirements for IBM XIV Storage Management GUI, XCLI, and XIVTop**
These requirements are necessary to install the latest IBM XIV Storage Management GUI, the IBM XIV online monitoring tool (XIVTop), and the IBM XIV command-line interface (XCLI).
- **GUI specifications**
The XIV GUI runs on either Direct mode or Manager mode. Additionally, the XIV GUI installed on a 64-bit OS has different specifications than 32-bit OS.

Requirements for IBM XIV Storage Management GUI, XCLI, and XIVTop

These requirements are necessary to install the latest IBM® XIV® Storage Management GUI, the IBM XIV online monitoring tool (XIVTop), and the IBM XIV command-line interface (XCLI).

The IBM XIV management tools include IBM XIV Storage Management GUI, IBM XIV command-line interface (XCLI), and IBM XIV online monitoring tool (XIVTop).

The following table lists the operating system and processor types that are supported by the IBM XIV management tools.

Table 1. Operating system and processor requirements

Operating system	Processor
(XCLI only) IBM AIX® 5.3, 6.1, or 7.1	IBM PowerPC®
(XCLI only) HP-UX 11i v3	Intel Itanium 64-bit
MAC OS X 10.6 or 10.7, with Java™ version 1.6 build 29 or later.	32-bit Dual Core or equivalent
Windows Server 2003, Windows Server 2008, Windows Vista, Windows 7	Intel Pentium Dual Core 32-bit, 64-bit, or equivalent
Red Hat Linux® version 5.0	x86 or x64
(XCLI only) Solaris 10	SPARC or x64

You must also satisfy the following requirements before you install the IBM XIV Storage Management GUI:

- Minimum of 100 MB of free space
- Minimum of 1 GB memory
- Minimum monitor resolution of 1024/768 (up to 1920/1200)
- Minimum monitor color resolution of 24/32 true color

Note: For more information, download the latest version of the XIV Management Tools Release Notes from www.ibm.com/support/knowledgecenter/STJTAG/com.ibm.help.xivgen3.doc/Gen3/xiv_gen3_kc_rn.html (www.ibm.com/support/knowledgecenter/STJTAG/com.ibm.help.xivgen3.doc/Gen3/xiv_gen3_kc_rn.html).

GUI specifications

The XIV GUI runs on either Direct mode or Manager mode. Additionally, the XIV GUI installed on a 64-bit OS has different specifications than 32-bit OS.

Direct mode

CPU

Dual core

Memory

Up to 25 systems

32-bit and 64-bit: 700MB

More than 25 systems

32-bit: 1.1GB

64-bit: 1.5GB

Max number of XIV systems monitored by a single XIV GUI in direct mode is:

81

Best practice: IBM XIV recommends to use the IBM® Hyper-Scale Manager and an XIV GUI in Manager mode if you have more than 10 systems.

Server mode

CPU

Dual core

Memory

500MB

Allowed latency between a single XIV GUI and the IBM Hyper-Scale Manager:

Max of 50ms

Supported OS for XIVGUI, XCLI, XIVTOP

Windows

- Windows Server 2003 (32 bit)
- Windows Server 2003 (64 bit)
- Windows Server 2008 (32 bit)
- Windows Server 2008 (64 bit)
- Windows Server 2012 (64 bit)
- Windows 7 (32 bit)
- Windows 7 (64 bit)
- Windows Vista (32 bit)
- Windows Vista (64 bit)
- Windows 8 (32 bit)
- Windows 8 (64 bit)

Make sure you install the relevant package (64-bit or 32-bit) according to your OS type.

Linux

- Linux Red Hat Enterprise 5
- Linux Red Hat 6.4 (32 bit)
- Linux Red Hat 6.4 (64 bit)

Mac

OSx 10.7

OSx 10.8

Additional supported OS for XCLI only

AIX

AIX 6

AIX 7

Oracle

Solaris 11 (SPARC)

Solaris 11 (INTEL)

HPUX

HPUX 11i v3 (IA64)

Planning for remote support and call home

This information describes what you must know to plan for remote support and call home connections.

- [Remote support connections](#)
XIV Remote Support Center is a management system used by authorized IBM service representatives to provide remote support and problem-determination assistance over a secure network interface, either before or after parts are called out by the internal diagnostic tests of the IBM XIV Storage System. IBM authorized personnel use remote support connections to perform real-time problem analysis and isolation.

- [Call home configuration](#)

Using call home, you can set up the XIV 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 XIV system to automatically send alerts directly to you. Call home cannot accept incoming communication, which means that the IBM Service Center cannot contact the XIV system using call home.

Remote support connections

XIV® Remote Support Center is a management system used by authorized IBM® service representatives to provide remote support and problem-determination assistance over a secure network interface, either before or after parts are called out by the internal diagnostic tests of the IBM XIV Storage System. IBM authorized personnel use remote support connections to perform real-time problem analysis and isolation.

Remote access uses one of the following two methods:

- A secure, high-speed, dedicated connection through a virtual private network (VPN) or management port on the XIV system. This is the preferred remote support method.
- A modem and dedicated, analog telephone line.
Note: The modem is no longer supported as part of the rack configuration for Indonesia and Brazil. For more information, contact your IBM planning representative.

The security features differ among these methods, including the direction of the connection.

- [Remote access through a dedicated network connection](#)

Remote access over a dedicated network connection is the preferred method for obtaining remote support. When using this method, the XIV system is connected to the XIV Remote Support Center through one of the virtual private network (VPN) ports or management ports on the patch panel.

- [Remote access by modem](#)

To use the modem for remote access, you must connect a dedicated, analog telephone line to the modem and give the telephone number to IBM. The telephone line must support dial-in; however, because the IBM XIV Storage System does not make outgoing calls by using the modem, the telephone line does not need to support dialing out.

Remote access through a dedicated network connection

Remote access over a dedicated network connection is the preferred method for obtaining remote support. When using this method, the XIV® system is connected to the XIV Remote Support Center through one of the virtual private network (VPN) ports or management ports on the patch panel.

The XIV Remote Support Center provides a high level of security for remotely supported XIV systems through encryption, authentication, authorization, auditing, and field-proven security components.


The XIV Remote Support Center has three components:

- The *XIV remote support client* is software that is installed on the XIV 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 XIV remote support client is controlled by using the IBM® XIV command-line interface (XCLI) and starts a connection, terminates a connection (due to timeout or customer request), and attempts to reconnect when the connection is terminated unexpectedly.
- The *front servers* serve as a hub at which the XIV systems and the remote-support back server connect. The front servers are located in an IBM DMZ and receive and maintain connections from the XIV 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 XIV systems 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, the XIV system must be able to initiate an outbound SSH connection to IBM. If the XIV system does not have direct access to the Internet (for example, due to a firewall), you can use the XIV Remote Support Proxy to facilitate the connection to IBM. For more information, see the *IBM XIV Storage System Remote Support Proxy Installation and User's Guide*.

The following figure illustrates the remote support components.

Figure 1. Remote support components



The XIV Remote Support Center 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 XIV Remote Support Center 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 XIV Remote Support Center.
2. The XIV Remote Support Center identifies the XIV system and marks it as connected.
3. The IBM service representative connects to the XIV Remote Support Center using SSH.
4. The XIV Remote Support Center authenticates the IBM service representative.
5. The IBM service representative is shown a list of currently connected XIV systems that correspond with the defined permissions, or the IBM service representative manually enters the serial number to view the XIV system.
6. The IBM service representative chooses the XIV system to support. Only permitted XIV 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 XIV system disconnects from the XIV Remote Support Center.

While a support session is in progress, the XIV system generates events and displays the system status on the IBM XIV Storage Management GUI. You can view the process of phasing out a module 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 XIV system next connects to the Remote Support Center.


Remote access by modem

To use the modem for remote access, you must connect a dedicated, analog telephone line to the modem and give the telephone number to IBM®. The telephone line must support dial-in; however, because the IBM XIV® Storage System does not make outgoing calls by using the modem, the telephone line does not need to support dialing out.

Note: The modem is no longer supported as part of the rack configuration for Indonesia and Brazil. For more information, contact your IBM planning representative. Ensure that you use the correct connector for your country to dial in. For example, in Germany and Luxembourg, use N-coded connectors.

[Figure 1](#) illustrates the remote support connectivity by using a modem.

Figure 1. Remote support connectivity by modem



Call home configuration

Using call home, you can set up the XIV® 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 XIV system to automatically send alerts directly to you. Call home cannot accept incoming communication, which means that the IBM Service Center cannot contact the XIV system using call home.

When certain events occur in the XIV 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 XIV system, IBM service personnel can perform service tasks, such as viewing error logs and problem logs or initiating trace and dump retrievals.

Call home notifications are sent through email from an SMTP server. You must have an SMTP email system available that the XIV system can use to send outgoing call home emails to IBM.

Important: It is a customer responsibility to configure the SMTP email system to enable the XIV 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, XIV 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 of the XIV system.

[Figure 1](#) shows call home routing.

Figure 1. Call home routing




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 the rear-door heat exchanger

If the optional rear-door heat exchanger (feature code 0082) was ordered, the customer must prepare the installation site before an IBM® service representative can install the rear-door heat exchanger on any of your XIV® systems.

- [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 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 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 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 rack inlet air, 20 kW rack load, 1530 cfm air through the rear-door heat exchanger)



Preparing your site to provide water to the 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](#)
Learn about the water specifications that are required 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

Learn about the water specifications that are required 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 Water delivery specifications for secondary loops), 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 - 6°C (39 - 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 $0.6 \text{ m} \times 0.6 \text{ m}$ ($2 \text{ ft} \times 2 \text{ ft}$) 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 $152.4 \text{ mm wide} \times 190.5 \text{ mm long} \pm 12.7 \text{ mm}$ ($6.0 \text{ in. wide} \times 7.5 \text{ in. long} \pm 0.5 \text{ in.}$) 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





Figure 2. Raised-floor hose management example 1: 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 $0.6 \text{ m} \times 0.6 \text{ m}$ ($2 \text{ ft} \times 2 \text{ ft}$) 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



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.

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. Miscellaneous secondary loop parts supplier information for customers in North America, Europe, Middle East, Africa, Asia Pacific

Supplier	Parts	Contact information
Vette Corporation	• Rear door heat exchangers (designed for non-IBM Enterprise racks) • Cooling distribution units • Hose kits • Water treatment • Chillers • Raised-floor grommets	Website: Vette Corp Email: datacom_facilities@vettecorp.com Phone: 877-248-3883 or 508-203-4690 Address: Vette Corp Datacom Facilities Division 201 Boston Post Road West Marlborough, MA 01752

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
Vette Corporation	- Installation of door and secondary loop items - Preventive maintenance	Website: Vette Corp Email: datacom_facilities@vettecorp.com Phone: 877-248-3883 or 508-203-4690 Address: Vette Corp Datacom Facilities Division 201 Boston Post Road West Marlborough, MA 01752

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 Eaton-Williams Group or Vette Corporation (see [Table 2](#)).

Table 3. Cooling distribution unit supplier information for customers in Europe

Supplier	Cooling distribution units	Contact information
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: Eaton-Williams Phone: (0) 1732 866055 Fax: (0) 1732 867937 Address: Eaton-Williams Group Ltd. Station Road Edenbridge Kent TN8 6EZ

The following illustration shows the Eaton-Williams cooling distribution unit features.

Figure 1. Eaton-Williams cooling distribution unit features



The cooling distribution unit specifications are described in the following table.

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)
Primary liquid connections	1 1/2 in. flex tail for sweat connection, top or bottom
Secondary liquid connections	3/4 in. quick connects, hydraulic ISO-B

Performance	Properties
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

Administering users

This information describes user roles and permissions and role-based authentication and access control.

- [User roles and permissions](#)
User roles allow specifying which roles are applied and the various applicable limits.
- [User authentication and access control](#)
IBM® XIV® Storage System features role-based authentication either natively or by using LDAP-based authentication.

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. Available user roles

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

User role	Permissions and limits	Typical usage
Storage administrator	The storage administrator has permission to all functions, except: - Maintenance of physical components or changing the status of physical components - Only the predefined administrator, named <i>admin</i>, can change the passwords of other users	Storage administrators are responsible for all administration functions.
Technician	The technician is limited to the following tasks: - Physical system maintenance - Phasing components in or out of service	Technicians maintain the physical components of the system. Only one predefined technician is specified per system. Technicians are IBM® XIV® Storage System technical support team members.

Notes:

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](#)
- [Predefined users](#)
- [User information](#)

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.

Predefined users

There are several predefined users configured on the IBM® XIV® Storage System. These users cannot be deleted.

Storage administrator

This user id provides the highest level of customer access to the system.

Predefined user name: `admin`

Default password: `adminadmin`

Technician

This user id is used only by IBM XIV Storage System service personnel.

Predefined user name: `technician`

Default password: Password is predefined and is used only by the IBM XIV Storage System technicians.

Note: Predefined users are always authenticated by the IBM XIV 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. for explanations of each role.

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: `~!@#$$%^&*()_+={}|;`
`<>?./\[]`.

E-mail

E-mail is used to notify specific users about events through e-mail messages. E-mail addresses must follow standard addressing procedures. E-mail is optional. Range: Any legal e-mail address.

Phone and area code

Phone numbers are used to send SMS messages to notify specific users about events. Phone numbers and area codes can be a maximum of 63 digits, hyphens (-) and periods (.) Range: Any legal telephone number; The default is N/A

User authentication and access control

IBM® XIV® 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 IBM XIV Storage System, regardless of the authentication mode.

- [Native authentication](#)
Native authentication is the default mode for authenticating users and user groups.
- [LDAP authentication](#)
Lightweight Directory Access Protocol (LDAP) support enables the IBM XIV Storage System to authenticate users through an LDAP repository.
- [LDAP authentication logic](#)
The LDAP authentication process consists of several key steps.

Native authentication

Native authentication is 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 IBM® XIV® Storage System to authenticate users through an LDAP repository.

When LDAP authentication is enabled, the username and password of a user accessing the IBM XIV Storage System (through CLI or GUI) are used by the IBM XIV system to login into a specified LDAP directory. Upon a successful login, the IBM XIV Storage System retrieves the user's IBM XIV group membership data stored in the LDAP directory, and uses that information to associate the user with an IBM XIV administrative role.

The IBM XIV group membership data is stored in a customer defined, pre-configured attribute on the LDAP directory. This attribute contains string values which are associated with IBM XIV administrative roles. These values might be LDAP Group Names, but this is not required by the IBM XIV Storage System. The values the attribute contains, and their association with IBM XIV administrative roles, are also defined by the customer.

Supported domains

The IBM XIV Storage System supports LDAP authentication of the following directories:

- 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 IBM XIV Storage System allows for more than one role for the Storage Administrator and the Read-Only roles.

The predefined XIV administrative IDs "admin" and "technician" are always authenticated by the IBM XIV storage system, whether or not LDAP authentication is enabled.

Responsibilities division between the LDAP directory and the storage system

Following are responsibilities and data maintained by the IBM XIV system and the LDAP directory:

LDAP directory

- Responsibilities - user authentication for IBM XIV users, and assignment of IBM XIV related group in LDAP.
- Maintains - Users, username, password, designated IBM XIV related LDAP groups associated with IBM XIV Storage System.

IBM XIV Storage System

- Responsibilities - Determination of appropriate user role by mapping LDAP group to an IBM XIV role, and enforcement of IBM XIV user system access.
- Maintains - Mapping of LDAP group to IBM XIV role.


LDAP authentication logic

The LDAP authentication process 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



User validation

During the login, the system validates the 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, base 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.

Managing

Refer to the following sections for information about managing IBM® XIV® Storage System instances using IBM Hyper-Scale Manager 5.x, IBM XIV Management Tools 4.8.x (including IBM Hyper-Scale Manager 1.9.x), and the IBM XIV command-line interface (XCLI).

- [Managing systems using IBM Hyper-Scale Manager 5.x](#)
For user information about IBM Hyper-Scale Manager, see the IBM Hyper-Scale Manager documentation.
- [Managing systems using IBM XIV Management Tools](#)
Information about managing IBM XIV Storage System instances using IBM XIV Management Tools 4.8.x and IBM Hyper-Scale Manager 1.9.x can be downloaded from the Publications and related information information.
- [Managing systems using the IBM XIV command-line interface \(XCLI\)](#)
For information about how to manage IBM XIV Storage System Gen3 instances using the IBM XIV command-line interface (XCLI), refer to the following section.

Managing systems using IBM Hyper-Scale Manager 5.x

For user information about IBM® Hyper-Scale Manager, see the IBM Hyper-Scale Manager documentation.

- [IBM Hyper-Scale Manager documentation](http://ibm.com/docs/en/hyper-scale-manager) (ibm.com/docs/en/hyper-scale-manager)

Managing systems using IBM XIV Management Tools

Information about managing IBM® XIV® Storage System instances using IBM XIV Management Tools 4.8.x and IBM Hyper-Scale Manager 1.9.x can be downloaded from the Publications and related information information.

Related reference

- [Publications and related information](#)

Managing systems using the IBM XIV command-line interface (XCLI)

For information about how to manage IBM® XIV® Storage System Gen3 instances using the IBM XIV command-line interface (XCLI), refer to the following section.

- [Command-line reference](#)

Operating

Use these operating procedures to start and stop the system and to install and update software.

- [Powering on and off the XIV system](#)
Use this information to power on and off the IBM® XIV® Storage System.
- [Installing and upgrading software](#)
This information describes prerequisites and procedures for installing management tools, considerations for updating system code, information about maintaining contact information, and information about upgrading IBM XIV Storage System code.

Powering on and off the XIV system

Use this information to power on and off the IBM® XIV® Storage System.

- [Powering on the XIV system, models 281x-114, 281x-116, 281x-214, and 281x-314, after shutdown](#)

Use this procedure to power on an IBM XIV Storage System, models 281x-114, 281x-116, 281x-214, and 281x-314, after either someone shut down the XIV system or after the XIV system automatically shut down.

- [Powering off XIV system models 281x-114, 281x-116, 281x-214, and 281x-314](#)

Use this procedure to power off the IBM XIV Storage System, models 281x-114, 281x-116, 281x-214, and 281x-314, by initiating a controlled shutdown.

Powering on the XIV system, models 281x-114, 281x-116, 281x-214, and 281x-314, after shutdown

Use this procedure to power on an IBM® XIV® Storage System, models 281x-114, 281x-116, 281x-214, and 281x-314, after either someone shut down the XIV system or after the XIV system automatically shut down.

About this task

(Models 281x-114, 281x-116, 281x-214, and 281x-314) The XIV system automatically performs an emergency shutdown when there is a power failure or when more than two uninterruptible power supplies failed. Emergency shutdowns are not performed by personnel. The XIV system determines when to carry out an emergency shutdown. Powering on the XIV system after an emergency shutdown uses the same procedure as when you recover the XIV system from a controlled shutdown.

If an external power loss occurs, the XIV system will go into emergency shutdown. When the external power returns, the XIV system will automatically power up. This will occur if the emergency shutdown was completed successfully (no interrupts of any kind). The system has a 6-minute timeout between the time the power resumed and the power-on. Should the system exhibit issues during the power-on operation (partial power on, boot to maintenance state, etc.), avoid all manual intervention and contact .

Attention: If you are recovering an XIV system after an emergency shutdown, determine the reasons for the shutdown. Correct the cause for the emergency shutdown before you power on the XIV system. If you cannot determine why the emergency shutdown occurred, contact an IBM service representative for assistance.

Procedure

Complete the following steps to power on the XIV system:

1. Ensure that external electrical sources are correctly supplying power to the XIV system.

Note: When both J1 and J2 power inlets are connected, J1 should be the default power connection.

If the power connection is forced to J2 (either by a bridge or by removing the J1 power connection) the power connection should return to J1 if the force mechanism is removed. (The force mechanism may be removed either by the bridge removal or if the power connection is returned to J1.

This behavior was released in automatic transfer switches (ATSs) manufactured ever since 2010/2011.

Prior to this, there was no default connection. Without the default connection, the connection was randomly selected and when you switched from J1 to J2, it stayed on J2, even after the force mechanism was removed.


[Table 1](#) specifies the exact model, part number, and dates for when the default behavior was introduced.

Table 1. ATS models and default J1 behavior

ATS Model	Part Number	Date
ATS60A	11S45W8823YM10NGBC3619	7 October 2010
IAC-30A	EU: 11S45W8827YM10NGAP0028	18 December 2010
	US: 11S45W8825YM10NGBN0817	16 November 2011
IAC-30A-3D	For this model there is no default, as both J1 and J2 are active, according to the following algorithm: IAC-30A-3D Normal mode Both line cords are usable. Two power outlets are fed from one line cord while the third power outlet is fed from the remaining line cord. IAC-30A-3D Abnormal mode a. Only one line cord is usable. • The power outlets are sourced from all three phases of this line cord. b. Both line cords are usable: • Only one outlet current exceeds 17A. This outlet is fed from one line cord only. • Two outlet currents exceed 17A. These two are fed from different line cords.	


2. From the rear of the rack, verify that the circuit breakers on the ATS are in the on (up) position.

Figure 1. Power sockets and circuit breakers on the side of the ATS




3. Verify that the three circuit breakers on each uninterruptible power supply are set to the on (up) position. If any of the circuit breakers are set to the off (down) position, do not proceed. Contact an IBM service representative for assistance.

Figure 2. Uninterruptible power-supply circuit breakers



4. From the front of the rack, press the Test button on each uninterruptible power supply in rapid succession, or simultaneously if possible, to power on the XIV system.

Figure 3. Front view of the uninterruptible power supply



5. Verify the system state

- [Verifying the system state](#)

Use this information to verify that the IBM XIV Software System components are operating correctly.

Verifying the system state

Use this information to verify that the IBM® XIV® Software System components are operating correctly.

Procedure

Complete the following steps to verify that the hardware components are operating correctly:

1. From the All Systems view in the IBM XIV Storage Management GUI, click the XIV system icon to display the system view.
2. Verify that the status pod in the lower-right corner of the window shows Full Redundancy.
If the state is not Full Redundancy, contact your IBM service representative.

3. If any component is in the Initializing state, wait 10 minutes, and check the status pod again.
If a component stays in the Initializing state (the component is displayed as yellow),
4. If any component is in the Ready state, complete the following steps from the IBM XIV Storage Management GUI:
 - a. Right-click the component, and click Phase In.
When you phase in a disk drive module (DDM), data module, or interface module, the component turns yellow, and the status on the lower right corner of the window shows Redistributing. Depending on the amount of data to be redistributed, this process might take some time.
 - b. Hover the cursor over the component, and verify that Status is OK.
5. If any component is in the Failed state (the component is displayed as red), complete the following steps:
 - a. If the reason cannot be determined, disconnect and then reconnect the power cable to the failed component, and then right-click the component, and click Test.
 - b. If the component is still failed, contact .
6. verify that there are no events with severities other than informational.
If there are other events, contact .

Powering off XIV system models 281x-114, 281x-116, 281x-214, and 281x-314

Use this procedure to power off the IBM® XIV® Storage System, models 281x-114, 281x-116, 281x-214, and 281x-314, by initiating a controlled shutdown.

Before you begin

Attention: Power off the XIV system only when necessary. When powered off, the XIV system is taken out of use and can no longer store data from host systems. Taking the XIV system out of use might put extra burden on all other XIV systems and might affect storage performance for the site.

About this task

(Models 281x-114, 281x-116, 281x-214, and 281x-314 only) The XIV system automatically performs an emergency shutdown when there is a power failure or when more than two uninterruptible power supplies fail. An emergency shutdown is not performed by personnel. The XIV system determines when to carry out an emergency shutdown.

When you power off the XIV system in a controlled fashion, the XIV system completes the following actions:

- Stops serving hosts
- Flushes all data from the cache, and places all the data onto the disks (also called *destaging data*)
- Turns itself off

Procedure

Complete the following steps to power off the XIV system:

1. Ensure that the XIV system is idle and ready to power off.
2. From the IBM XIV Storage Management GUI, click the user ID in the upper-right corner, and log in using the storage administrator user ID and password.
3. Click the XIV system to be powered off.
4. Click System > System Settings > Shutdown System.
5. Click OK.
6. If your XIV system is to be powered off for more than 4 hours, contact an IBM service representative to disconnect the uninterruptible power supply batteries to preserve the battery charge.

Installing and upgrading software

This information describes prerequisites and procedures for installing management tools, considerations for updating system code, information about maintaining contact information, and information about upgrading IBM® XIV® Storage System code.

- **Non-disruptive code load**
Non disruptive code load (a.k.a. Hot Upgrade) enables the IBM XIV Storage System to upgrade its software from a current version to a newer version without disrupting application service.

Non-disruptive code load

Non disruptive code load (a.k.a. Hot Upgrade) enables the IBM® XIV® Storage System to upgrade its software from a current version to a newer version without disrupting application service.

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 damaged. The upgrade requires that neither data migration nor rebuild processes are run, and that all internal network paths are active.

During the non disruptive code load process there is a point in time dubbed the 'upgrade-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 Non-Disruptive Code Load process is not reversible.

Following are notable characteristics of the 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 online to the application/Host.
The duration of the upgrade process is affected by the following factors:

- The upgrade process requires that you reduce all IOs. If there are a lot of IOs in the system, or there are slow disks, the system might not be able to stop the IOs fast enough, so it will restart them and try again after a short while, taking into consideration some retries.
- The upgrade process installs a valid version of the software and then retains its local configuration. This process might take a considerable amount of time, depending on the future changes in the structure of the configuration.

Prerequisites and constraints

- The process cannot run if a data migration process or a rebuild process is active. An attempt to start the upgrade process when either a data migration or a rebuild process is active will fail.
- Generally, everything that happens after the point-of-no-return is treated as if it happened after the upgrade is over.
- 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.
- Prior to the point-of-no-return, a manual abort of the upgrade is available.

Effect on mirroring

Mirrors are automatically deactivated before the upgrade, and reactivated after it is over.

Effect on management operations

During the Non-Disruptive Code Load 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 will be aborted automatically.

Handling module or disk failure during the upgrade

If the failure occurs before the point-of-no-return, it will abort the upgrade. If it happens after that point, the failure is treated as if it happened after the upgrade is over.

Handling power failure during the upgrade

As for power failure before the point-of-no-return - power is being monitored during the time the system prepares for the upgrade (before the point-of-no-return). If a power failure is detected, the upgrade will be aborted and the power failure will be taken care of by the old version.

- [Preparing for code load](#)

The IBM XIV Storage System upgrades the system code without disconnecting active hosts or stopping I/O operations.

Preparing for code load

The IBM® XIV® Storage System upgrades the system code without disconnecting active hosts or stopping I/O operations.

Non-disruptive code load is an operation performed by IBM XIV personnel only. This guide includes mandatory and best practice procedures applicable to normal operations and upgrades.

Preparing for the upgrade

Before the code load, fulfill the following prerequisites by verifying that:

1. The multi-pathing feature (provided by the operating system) is working on the host.
2. There are paths from the host to at least two different interface modules on IBM XIV.
3. There is no more than a single initiator in each zone (SAN Volume Controller attached to IBM XIV is an exception).
4. The host was attached to IBM XIV using the xiv_attach utility.
 - This is mandatory.
 - This applies to both installable HAK and portable HAK.
 - Exceptions to this prerequisite are supported platforms, for which no HAK is available (for example, VMware or Linux on Power systems).
5. The minimal version of the "IBM XIV Host Attachment Kit for Windows" is 1.5.3. This version prevents Windows hosts from a potential loss of access.
6. In a case of IBM XIV uses 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.
7. 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 IBM XIV Storage System is not supported.

Be aware of the following:

1. Co-existence with other multi-pathing 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. The mirroring is automatically suspended and resumed for a short while during the upgrade.
4. Mirroring from 10.2.4.x to 10.2.1.x or older versions is not supported.
5. There are special considerations where MS Geo Cluster is involved. Contact IBM support for more details.

Recommended practices:

1. It is highly recommended to have the latest XIV Host Attachment Kit installed on the host. Each release of a HAK fixes known issues on older versions.
 - Being on an older level means being exposed to problems already found and fixed.
2. It is recommended to follow the zoning best practices of IBM XIV as described in the Redbook and in the XIV Host Attachment Guide.
3. Best practice is to follow the OS Provider recommendations regarding service packs and storage-related Hot Fixes. These fixes are released from time to time by the OS provider and are outside of IBM XIV control. Some fixes are listed on the release notes of the latest available HAK and have to be applied before the upgrade.
4. It is recommended to keep your host system with up-to-date BIOS and HBA drivers.

5. It is recommended to perform the upgrade on time during which the workload is relatively low.

Availability of TA support:

1. If entitled with a TA service for your IBM XIV Storage System, contact your assigned Technical Advisor when planning for a code upgrade.

Remote support

Refer to the following sections for information about remote support available for IBM® XIV® Storage System.

- [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 preformed.

- **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.
- **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 XIV® Remote Support Proxy utility to enable the connection. Remote support minimizes the time it takes to diagnose and remedy storage system operational issues.

Note: No data can be accessed by IBM Support when the remote support option is used.

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 exist.

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 exist.

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=129.33.206.139:443,204.146.30.139:443,170.225.126.11:443,170.225.126.12:443,170.225.127.11:443,170.225.127.12:443
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] IPAddr
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.

Reference

Refer to the following sections for information about how to use the IBM® XIV® command-line interface (XCLI), IBM Hyper-Scale Manager REST Application Programming Interface (API), and the XIV Open Application Program Interface (API).

- [Command-line reference](#)
The following information describes how to use the IBM XIV command-line interface (XCLI).
- [REST API reference](#)
For information about how to use the IBM Hyper-Scale Manager REST Application Programming Interface (API), see the IBM Hyper-Scale Manager documentation.
- [XIV Open API](#)
The following information describes how to use the IBM XIV Software System Open Application Programming Interface (API).

Command-line reference

The following information describes how to use the IBM® XIV® command-line interface (XCLI).

- [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.
- [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 volume 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.
- [Grouped pool management commands](#)
- [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.
- [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.
- [Compression commands](#)
- [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.
- [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.
- [Metadata-service commands](#)
This section describes the command-line interface (CLI) for handling the metadata service.
- [Encryption enablement and support commands](#)
This section describes the command-line interface (CLI) for encryption configuration.
- [xmirror commands](#)
This section describes the command-line interface (CLI) for xmirror creation and management.
- [Security configuration commands](#)
- [Events](#)
This section contains detailed information on CLI events, including their severity and descriptions.
- [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. The XCLI is available on Microsoft Windows, Linux® and other operating systems.
- [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. The XCLI is available on Microsoft Windows, Linux® and other operating systems.

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 notes 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 or clusters](#)
Use the **mapping_list** command to list the mapping of volumes to a specified host or cluster.
- [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.
- [Setting 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 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.
- [Creating an ALU LUN](#)
Use the `alu_create` command to create a new ALU and map it to a given host.
- [Deleting an ALU LUN](#)
Use the `alu_delete` command to delete an existing ALU.
- [Listing ALUs](#)
Use the `alu_list` command to list all ALU LUNs in the storage system.
- [Binding volumes to an ALU](#)
Use the `alu_bind` command to attach a list of existing volumes to an ALU.
- [Unbinding volumes from an ALU](#)
Use the `alu_unbind` command to detach a list of volumes from an ALU.
- [Unbinding all volumes from an ALU](#)
Use the `alu_unbind_all` command to detach all SLUs from the specified ALU.
- [Listing all volumes bound to an ALU](#)
Use the `alu_bind_list` command to list all volumes bound to an ALU.
- [Listing all ALUs and hosts to which the volume is bound](#)
Use the `vol_bind_list` command to list all ALUs and hosts to which the volume is bound.

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 overridden 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**
Host name does not exist
- **CLUSTER_BAD_NAME**
Cluster name does not exist
- **HOST_BELONGS_TO_ANOTHER_CLUSTER**
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**
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
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- **CLUSTER_NAME_EXISTS**
Cluster name already exists
- **MAX_CLUSTERS_REACHED**
Maximum number of clusters already defined
- **DOMAIN_DOESNT_EXIST**
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'* has hosts in it. Are you sure you want to delete it?

Return codes

- CLUSTER_BAD_NAME
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

Field ID	Field output	Default position
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
Host name does not exist
- CLUSTER_BAD_NAME
Cluster name does not exist
- HOST_NOT_IN_CLUSTER
Host is not part of 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
Cluster name does not exist
- CLUSTER_NAME_EXISTS
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 >
```

Parameters

Name	Type	Description	Mandatory
host	Object name	The host name.	Y
fcaddress	N/A	FC address of the added port.	N
iscsi_name	iSCSI initiator name	iSCSI initiator name of the newly added port.	N

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
Host name does not exist
- HOST_PORT_EXISTS
Host with this port ID already defined
- ISCSI_HOST_ILLEGAL_PORT_NAME
Port name for iSCSI Host is illegal

Troubleshooting: Port names for iSCSI Hosts must contain only printable characters.

- **MAX_PORTS_REACHED**
Maximum number of ports already defined in the system
- **TARGET_PORT_BAD_ADDRESS**
Remote port address is illegal or does not belong to the remote target
- **PORT_EXISTS**
Port is already defined
- **OLVM_LINK_IS_NOT_UP**
IBM Hyper-Scale Mobility link is not up. The mapping list cannot be updated.
- **REMOTE_MAX_VIRTUAL_HOSTS_REACHED**
Maximum number of remote virtual hosts already defined

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 executed 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**
Host name already exists
- **MAX_HOSTS_REACHED**
Maximum number of hosts already defined
- **CLUSTER_BAD_NAME**
Cluster name does not exist
- **DOMAIN_DOESNT_EXIST**
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
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**
Host name does not exist
 - **HOST_IS_ALU_MAPPED**
Host that is mapped to an ALU cannot be deleted
- Troubleshooting:** Delete the ALU that maps the host

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
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=tlb_host_pro125_fc0
```

Output:

Host	Type	Port name
tlb_host_pro125_fc0	FC	100000062B125CD0

Field ID	Field output	Default position
host	Host	1
type	Type	2
port_name	Port Name	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

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:

```
xcli.py 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**
Host name does not exist
- **PORT_DOES_NOT_BELONG_TO_HOST**
Port ID belongs to another host
- **HOST_PORT_DOES_NOT_EXIST**
Port ID is not defined
- **ISCSI_HOST_ILLEGAL_PORT_NAME**
Port name for iSCSI Host is illegal

Troubleshooting: Port names for iSCSI Hosts must contain only printable characters.

- **OLVM_LINK_IS_NOT_UP**
IBM Hyper-Scale Mobility link is not up. The mapping list cannot be updated.
- **TARGET_PORT_BAD_ADDRESS**
Remote port address is illegal or does not belong to the remote target
- **HOST_PORT_EXISTS**
Host with this port ID already defined
- **MAX_PORTS_REACHED**

Maximum number of ports already defined in the system

- **PORT_EXISTS**
Port is already defined
- **REMOTE_MAX_VIRTUAL_HOSTS_REACHED**
Maximum number of remote virtual hosts already defined

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
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- **HOST_BAD_NAME**
Host name does not exist
- **HOST_NAME_EXISTS**
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 executed 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`
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 wrong size - 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
Host name does not exist
- HOST_BELONGS_TO_CLUSTER
Host is part of a cluster
- CLUSTER_BAD_NAME
Cluster name does not exist
- VOLUME_BAD_NAME
Volume name does not exist
- SNAPSHOT_IS_INTERNAL
Internal snapshots cannot be mapped, modified or deleted.
- VOLUME_ALREADY_ASSIGNED
Mapping conflict: volume is already assigned
- LUN_ALREADY_IN_USE
Mapping conflict: LUN is already in use
- EXT_LUN_ILLEGAL
LUN is out of range or does not exist
- VOLUME_HAS_HOST_SPECIFIC_MAPPING
Specified Volume is currently mapped to another LUN in a host-specific mapping
- LUN_HAS_HOST_SPECIFIC_MAPPING
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
Port name for iSCSI Host is illegal

Troubleshooting: Port names for iSCSI Hosts must contain only printable characters.
- MAX_PORTS_REACHED
Maximum number of ports already defined in the system
- OLVM_LINK_IS_NOT_UP
IBM Hyper-Scale Mobility link is not up. The mapping list cannot be updated.
- HOST_PORT_EXISTS

Host with this port ID 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
Maximum number of remote virtual hosts already defined
- CAN_NOT_MAP_SLAVE_COMPRESSED_VOLUME
Slave compressed volume cannot be mapped
- VOLUME_HAS_INACTIVE_DATA_MIGRATION
Cannot map a volume that has an inactive data migration

Listing the mapping of volumes to hosts or clusters

Use the **mapping_list** command to list the mapping of volumes to a specified host or cluster.

```
mapping_list [ host=HostName | cluster=ClusterName ] [ domain=DomainList ]
```

Parameters

Name	Type	Description	Mandatory	Default
host	Object name	Host name.	N	N/A
cluster	Object name	Cluster name.	N	N/A
domain	N/A	List of hosts, clusters or domains to show mapping from. To define more than one host, cluster or 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		
serial	Serial Number	6		
locked	Locked	7		
host	Host	8		

Example:

```
mapping_list host=demo_host_1,demo_host_fc10000006072d0190
```

Output:

LUN	Volume	Size	Master	Serial Number	Locked	Host
0	vol-870834-0003	137		3	no	demo_host_fc10000006072d0190
1	vol-870834-0004	137		4	no	demo_host_fc10000006072d0190
2	vol-870834-0005	137		5	no	demo_host_fc10000006072d0190
3	vol-870834-0006	137		6	no	demo_host_fc10000006072d0190
4	vol-870834-0007	34		7	yes	demo_host_fc10000006072d0190
5	vol-870834-0008	34		8	no	demo_host_fc10000006072d0190
6	vol-870834-0010	34		10	no	demo_host_fc10000006072d0190
7	vol-870834-0009	34		9	no	demo_host_fc10000006072d0190
8	vol-870834-0011	34		11	no	demo_host_fc10000006072d0190
9	vol-870837-0004	17		12	no	demo_host_fc10000006072d0190
10	vol-870837-0006	17		13	no	demo_host_fc10000006072d0190
11	vol-870837-0022	17		27	no	demo_host_fc10000006072d0190
12	vol-870837-0024	17		28	no	demo_host_fc10000006072d0190
13	vol-870837-0027	68		31	no	demo_host_fc10000006072d0190
14	vol-870837-0028	86		32	no	demo_host_fc10000006072d0190

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**
Host name does not exist
- **CLUSTER_BAD_NAME**
Cluster name does not exist
- **TOO_MANY_MAPPINGS**
There are too many mappings to display
- **DOMAIN_DOESNT_EXIST**
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>
```

Parameters

Name	Type	Description	Mandatory
host	Object name	Host name.	N
cluster	Object name	Cluster name.	N
type	Enumeration	Special map type.	Y

The supported special types are **hpux** and **zvm**. The type must be specified for hosts or clusters that run on the HP-UX operating system. All other operating systems do not require a special type.

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.

Example:

```
special_type_set host=tlib_host_pro26_fc0 type=zvm
```

Output:

Command executed 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**
Host name does not exist
- **HOST_BELONGS_TO_CLUSTER**
Host is part of a cluster
- **CLUSTER_BAD_NAME**
Cluster 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.

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

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
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
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 ]
```

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

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 again the command.

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

- HOST_BAD_NAME
Host name does not exist
- HOST_BELONGS_TO_CLUSTER

Host is part of a cluster

- **CLUSTER_BAD_NAME**
Cluster name does not exist
- **VOLUME_BAD_NAME**
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**
IBM Hyper-Scale Mobility link is not up. The mapping list cannot be updated.
- **ISCSI_HOST_ILLEGAL_PORT_NAME**
Port name for iSCSI Host is illegal

Troubleshooting: Port names for iSCSI Hosts must contain only printable characters.

- **MAX_PORTS_REACHED**
Maximum number of ports already defined in the system
- **HOST_PORT_EXISTS**
Host with this port ID already defined
- **VOLUME_NOT_IDLE**
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**
Maximum number of remote virtual hosts already defined

Setting 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

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

The performance class name must be unique. Up to 1000 classes can be created.

Example:

```
perf_class_create perf_class=p1
```

Output:

```
Command executed 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
Performance Class already exists.
- MAX_PERF_CLASSES_REACHED
Maximum number of Performance Class is already defined.

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 executed 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_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**
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 executed 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**
Performance Class already exists.
- **PERF_CLASS_BAD_NAME**
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.
ID	Name	Default Position		
name	Performance class	1		
max_iops	Max IOPS rate (IOPS)	3		
max_bw	Max bandwidth rate (MB/s)	4		

Example:

```
perf_class_list
```

Output:

```
Performance class  Max IO rate (IOPS)  Max BW rate (MB/s)
perf1              0                  0
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
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 executed 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
Performance Class does not exist
- HOST_BAD_NAME
Host name does not exist
- HOST_ALREADY_IN_PERF_CLASS
Host *host* already in Performance Class *Performance Class*.
- PERF_CLASS_ASSOCIATED_WITH_POOLS_OR_DOMAINS
Performance Class *Performance Class* is already in use by pool or domain.

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 executed 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
Host name does not exist
- PERF_CLASS_DOES_NOT_CONTAIN_ANY_HOSTS
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 executed 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**
Performance Class does not exist
- **POOL_DOES_NOT_EXIST**
Storage Pool does not exist
- **POOL_ALREADY_IN_PERF_CLASS**
Pool *pool name* already in Performance Class *Performance Class*.
- **PERF_CLASS_ASSOCIATED_WITH_HOSTS**
Performance Class *Performance Class* is already in use by host.

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 executed 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**
Storage Pool does not exist
- **POOL_NOT_CONNECTED_TO_ANY_PERF_CLASS**
Pool 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

Name	Type	Description	Mandatory
perf_class	Object name	The name of a performance class.	Y

Example:

```
perf_class_add_domain perf_class=perf1 domain=d1
```

Output:

```
Command executed 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
Domain does not exist.
- PERF_CLASS_BAD_NAME
Performance Class does not exist
- DOMAIN_ALREADY_IN_PERF_CLASS
Domain *domain name* already in Performance Class *Performance Class*.
- PERF_CLASS_ASSOCIATED_WITH_HOSTS
Performance Class *Performance Class* is already in use by host.

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 executed 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
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 (IOPS). The max 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 in bandwidth per interface (MB/sec). The max setting allowed is 10,000. If zero is specified, the bandwidth rate will not be limited.	N	Keep unchanged.

This command sets the rate for a performance class. The specified rate is applied for each interface module. Either max_io_rate, max_bw_rate or both must be set.

Example:

```
perf_class_set_rate perf_class=p1 max_io_rate=1000
```

Output:

```
Command executed 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
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
-----
```

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
```

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 executed 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.
- **HOST_BAD_NAME**
Host name does not exist
- **MAX_HOST_PROFILES_REACHED**
Maximum number of host profiles already defined
- **HOST_PROFILER_DISABLED**
Host profiler disabled

Troubleshooting: Try again after the minimal update interval time

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**
Host name does not exist
- **HOST_PROFILE_DOES_NOT_EXIST**
No profile 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

Creating an ALU LUN

Use the **alu_create** command to create a new ALU and map it to a given host.

```
alu_create alu=ALUName host=HostName lun=LUN
```

Parameters

Name	Type	Description	Mandatory
host	Object name	The name of the host the ALU is mapped to.	Y
lun	Integer	The LUN of the ALU in the host (must be between 513 and 754).	Y
alu	Object name	The ALU name	Y

The ALU represents the host for VM volume operations. ALU names are unique per system, even when multi-tenancy is supported.

Example:

```
alu_create alu=alu_name host=myHost lun=600
```

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

- HOST_BAD_NAME
Host name does not exist
- ALU_ALREADY_EXISTS
ALU name already exists
- MAX_ALUS_REACHED
Maximum number of ALU already defined
- LUN_ALREADY_IN_USE
Mapping conflict: LUN is already in use

Deleting an ALU LUN

Use the **alu_delete** command to delete an existing ALU.

```
alu_delete alu=ALUName
```

Parameters

Name	Type	Description	Mandatory
alu	Object name	The name of the ALU to be deleted.	Y

To be suitable for deletion, the ALU must not have any (2nd level) LUNs associated with it.

Example:

```
alu_delete alu=alu1
```

Access control

User Category	Permission
Storage administrator	Disallowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_DELETE_ALU
Are you sure you want to delete ALU *ALU*?

Return codes

- ALU_HAS_BOUND_SLU_CANNOT_DELETE
Volumes are bounded to this ALU

Troubleshooting: Unbind all volumes from ALU

- ALU_BAD_NAME
ALU name does not exist

Listing ALUs

Use the **alu_list** command to list all ALU LUNs in the storage system.

```
alu_list [ host=HostName ] [ domain=DomainName ]
```

Parameters

Name	Type	Description	Mandatory	Default
host	Object name	The host name.	N	All hosts.
domain	Object name	The domain name.	N	All Domains

This command returns a list of triplets of host, ALU, and lun_id, optionally filtered by host.

A host name can be specified to list either the specific ALU LUNs exposed by this host, or all ALU LUNs.

The list contains the following comma-separated information:

- ALU name
- Host name
- LUN number

Example:

```
alu_list host
```

Output:

```
Name      Host Name  Lun
-----
alu1      host1      754
```

ID	Name	Default Position
name	Name	1
host	Host Name	2
lun	Lun	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

Binding volumes to an ALU

Use the **alu_bind** command to attach a list of existing volumes to an ALU.

```
alu_bind alu=ALUName vols=<vol1[,vol2]...>
```

Parameters

Name	Type	Description	Mandatory
vols	Object name	List of all the volumes to bind to the ALU. Up to 256 SLUs can be listed.	Y
alu	Object name	The name of the ALU the volume will be bound to.	Y

If a volume in the list is already attached to the ALU, no change will be applied. A volume can be bound to more than one ALU at the same time.

This command attempts to bind as many volumes as possible, and does not stop on error (that is, if the binding fails for any volume). The command returns identical results in SUCCESS and SUCCESS_PARTIAL, so that the calling application can determine the slun_id of each bound volume.

Up to 256 SLUs can be listed.

Example:

```
alu_bind alu vols
```

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

- VOLUME_BAD_NAME
Volume name does not exist
- ALU_BAD_NAME
ALU name does not exist
- ALL_VOLUMES_MUST_BE_FROM_SAME_DOMAIN
All the volumes in a single command must belong to a single domain.
- VOLUME_HAS_TRANSFORM
Operation is not supported for volumes which are compressed/decompressed

Unbinding volumes from an ALU

Use the **alu_unbind** command to detach a list of volumes from an ALU.

```
alu_unbind alu=ALUName vols=<vol1[,vol2]...>
```

Parameters

Name	Type	Description	Mandatory
vols	Object name	List of all the volumes to unbind from the ALU. Up to 256 SLUs can be listed.	Y
alu	Object name	The name of the ALU the volume will be unbound from.	Y

If a volume in the list is already detached from the ALU, the command completes successfully, but has no effect.

Example:

```
alu_unbind alu vols
```

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

- VOLUME_BAD_NAME
Volume name does not exist
- ALU_BAD_NAME
ALU name does not exist
- ALL_VOLUMES_MUST_BE_FROM_SAME_DOMAIN
All the volumes in a single command must belong to a single domain.

Unbinding all volumes from an ALU

Use the **alu_unbind_all** command to detach all SLUs from the specified ALU.

```
alu_unbind_all alu=ALUName
```

Parameters

Name	Type	Description	Mandatory
alu	Object name	The name of the ALU the volumes will be unbound from.	Y

Example:

```
alu_unbind_all alu
```

Access control

User Category	Permission
Storage administrator	Disallowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_UNBIND_ALL
Are you sure you want to unbind all volumes for ALU *ALU*?

Return codes

- ALU_BAD_NAME
ALU name does not exist

Listing all volumes bound to an ALU

Use the **alu_bind_list** command to list all volumes bound to an ALU.

```
alu_bind_list alu=ALUName
```

Parameters

Name	Type	Description	Mandatory
alu	Object name	The name of the ALU to retrieve information about.	Y

This command returns a list of couples of (slu_name, slu_lun), where ALU and SLU are bound. The format of the slu_lun is D2YYYYYY0000 (hexadecimal), where YYYYYY is the 24 bits that uniquely identify a bound volume.

Example:

```
alu_bind_list alu
```

Output:

```
SLU Name    SLU LUN
-----
v1          0000D20000030000
```

ID	Name	Default Position
slu_name	SLU Name	1
slu_lun	SLU LUN	2

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed

User Category	Permission
Application administrator	Allowed
Security administrator	Allowed
Read-only users	Allowed
Technicians	Allowed

Return codes

- ALU_BAD_NAME
ALU name does not exist

Listing all ALUs and hosts to which the volume is bound

Use the **vol_bind_list** command to list all ALUs and hosts to which the volume is bound.

```
vol_bind_list vol=VolName
```

Parameters

Name	Type	Description	Mandatory
vol	Object name	Volume name.	Y

The list contains slu_lun, alu name, alu_lun, and host. The format of the slu_lun is D2YYYYYY0000 (hexadecimal), where YYYYYY is the 24 bits, that uniquely identify a bound volume.

Example:

```
vol_bind_list vol=alu_vol
```

Output:

```
SLU Name    SLU LUN
-----
v1          0000D20000030000
```

Field ID	Field output	Default position
slu_lun	SLU LUN	1
alu	ALU	2
host	HOST	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

- VOLUME_BAD_NAME
Volume name does not exist

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.
- [Finding a volume by its internal index](#)
Use the **vol_by_id** command to print the volume name according to its internal volume index.
- [Copying volumes](#)
Use the **vol_copy** command to copy a source volume onto a target volume.
- [Copying volume differences](#)
Use the **vol_diff_copy** command to copy the difference between a previously copied volume and its base volume onto a target volume.
- [Creating a volume](#)
Use the **vol_create** command to create a new volume.
- [Creating multiple volumes](#)
Use the **vol_create_many** command to create new volumes.
- [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.
- [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.
- [Compressing volumes](#)
Use the **vol_transform_create** command to compress a regular volume or decompress a compressed volume.
- [Managing volume compression and decompression](#)
Use the **vol_transform_update** command to control volume compression and decompression transformation requests.
- [Deleting a compression or decompression request](#)
Use the **vol_transform_delete** command to delete a volume compression or decompression request from the queue.
- [Listing the volume compression activity](#)
Use the **vol_transform_list** command to list the state of all volume transformation requests.
- [Setting a volume's external identifier](#)
Use the **vol_set_external_id** command to set a volume's external identifier.
- [Clearing a volume's external identifier](#)
Use the **vol_clear_external_id** command to clear a volume's external identifier.
- [Displaying the allocated partitions of a volume](#)
Use the **vol_allocation_bitmap** command to display the allocated partitions of a volume.
- [Displaying the difference of allocated partitions between two volumes](#)
Use the **vol_diff** command to display the allocated partitions' difference between two volumes.
- [Triggering compression estimation for a system volume](#)
Use the **vol_comprestimate** command to request an estimate of a volume's compression ratio.
- [Removing a volume from the compression estimation queue](#)
Use the **vol_comprestimate_stop** command to remove the specified volume from the compression estimation queue.
- [Displaying the volume's estimated compression ratio](#)
Use the **vol_comprestimate_list** to display the volume's compression estimates.

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
Volume name does not exist

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
Technicians	Disallowed

Return codes

- VOLUME_BAD_NAME
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
Volume name does not exist

Finding a volume by its internal index

Use the **vol_by_id** command to print the volume name according to its internal volume index.

```
vol_by_id id=n
```

Parameters

Name	Type	Description	Mandatory
id	Positive integer	SCSI volume ID.	Y

Example:

```
vol_by_id id=59
```

Output:

```
Name      Size (GB) Master Name Consistency Group Pool Creator Used Capacity(GB)
volume_1  51                                0
```

Field ID	Field output	Default position
name	Name	1
size	Size (GB)	2
size_MiB	Size (MiB)	N/A
master_name	Master Name	3
cg_name	Consistency Group	4
pool_name	Pool	5
creator	Creator	6
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
serial	Serial Number	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

Field ID	Field output	Default position
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
used_capacity	Used Capacity (GB)	9
used_capacity_MiB	Used Capacity (MiB)	N/A
capacity_used_by_snapshots_MiB	Capacity Used by Snapshots (MiB)	N/A
short_lived_io	Short Live IO	N/A
enabled_VAAI_cmd	VAAI commands enabled bitmask	N/A
user_disabled_VAAI	VAAI disabled by user	N/A
snapshot_format	Snapshot Format	N/A
compressed	Compressed	7
ratio	Compression Ratio (%)	8
saving	Compression Saving (GB)	10
saving_MiB	Compression Saving (MiB)	N/A
ssd_caching	SSD Caching State	N/A
use_ssd_caching_default	Use SSD Caching Default State	N/A
unmap_support	Unmap Support	N/A
managed	Managed	N/A
enable_unmap	unmap enabled	N/A
user_disabled_unmap	unmap disabled by user	N/A
marked	Marked	N/A
online	Online	N/A
inflated	Inflated	N/A
metadata_mismatch	MetaData Mismatch	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Operations administrator	Disallowed
Host side accelerator client	Disallowed

Return codes

- VOLUME_BAD_SERIAL
Volume with requested SCSI serial number 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 than 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 executed 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_HARD_SPACE
No space to allocate for volume's current usage
- NOT_ENOUGH_SPACE
No space to allocate volume
- SOURCE_VOLUME_BAD_NAME
Source volume name does not exist
- SOURCE_VOLUME_DATA_MIGRATION_UNSYNCHRONIZED
Data Migration has not completed to source volume
- TARGET_VOLUME_BAD_NAME
Target volume name does not exist
- TARGET_VOLUME_LOCKED
Target volume is locked
- TARGET_VOLUME_HAS_MIRROR
Mirror is defined for target volume
- TARGET_VOLUME_HAS_DATA_MIGRATION
Data Migration is defined for target volume
- VOLUME_IS_SNAPSHOT
Operation is not permitted on snapshots
- VOLUME_IDENTICAL
Source and target are the same volume
- VOLUME_HAS_SNAPSHOTS
Operation is not permitted on volume with snapshots
- VOLUME_IS_NOT_CONSISTENT_SLAVE
Operation not allowed on slave volume that is not consistent.

- **TARGET_VOLUME_NOT_FORMATTED**
Target volume is not formatted
- **SNAPSHOT_IS_FORMATTED**
Snapshot is formatted
- **VOLUME_TOO_BIG_TO_COPY**
Volume is too large to be copied
- **VOLUME_HAS_TRANSFORM**
Operation is not supported for volumes which are compressed/decompressed
- **TARGET_VOLUME_HAS_OLVM**
This target volume is part of an IBM Hyper-Scale Mobility relation.
- **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.
- **SOURCE_VOLUME_COMPRESSED_TARGET_UNCOMPRESSED**
Source volume is compressed but target volume is uncompressed.
- **SOURCE_VOLUME_UNCOMPRESSED_TARGET_COMPRESSED**
Source volume is uncompressed but target volume is compressed.
- **COMPRESSED_CAPACITY_LIMIT_REACHED**
Total capacity of compressed volumes exceeds the maximum limit
- **COMPRESSION_VOLUME_IS_OFFLINE**
Operation cannot be completed because the requested compressed volume is offline

Copying volume differences

Use the **vol_diff_copy** command to copy the difference between a previously copied volume and its base volume onto a target volume.

```
vol_diff_copy vol_src=VolName vol_base=VolName vol_trg=VolName
```

Parameters

Name	Type	Description	Mandatory
vol_src	Object name	Name of the source volume from which the different data from the base volume is to be taken.	Y
vol_base	Object name	Name of the base volume to which the data in the source volume is compared.	Y
vol_trg	Object name	Name of the target volume to which the data is to be copied.	Y

This command copies the difference between a source volume and its base volume onto a target volume that belongs to a different volume tree.

All volume-to-host mappings of the three volumes remain intact during this process. The target volume retains all of its properties, including its name, ID, lock state, creation time and all other attributes.

The command fails in the following cases:

- The three volumes do not have the same size.
- The source and base volumes do not belong to the same volume tree.
- The source volume is formatted.
- The target volume belongs to the source volume tree.
- The target volume is not a volume tree root.
- 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 definition, and it is currently inconsistent (due to either a re-synchronization or an initialization process).

Example:

```
vol_diff_copy vol_src=DBVolume vol_base=DBVolumeSnap vol_trg=DBVolumeCopy
```

Output:

```
Command executed successfully
```

Access control

User Category	Permission
Storage administrator	Disallowed
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_COPY_DIFF_VOLUME
Are you sure you want to copy the contents of volume *source Volume* from base *base Volume* to volume *target Volume*?

Return codes

- SOURCE_VOLUME_BAD_NAME
Source volume name does not exist
- SOURCE_VOLUME_DATA_MIGRATION_UNSYNCHRONIZED
Data Migration has not completed to source volume
- TARGET_VOLUME_BAD_NAME
Target volume name does not exist
- TARGET_VOLUME_LOCKED
Target volume is locked
- TARGET_VOLUME_HAS_MIRROR
Mirror is defined for target volume
- TARGET_VOLUME_HAS_DATA_MIGRATION
Data Migration is defined for target volume
- TARGET_VOLUME_IS_NOT_ROOT
Target volume must be master and not after restore
- BASE_VOLUME_BAD_NAME
Base volume name does not exist
- BASE_VOLUME_INVALID
Base and source volumes must be different and belong to the same master volume
- VOLUME_IS_NOT_CONSISTENT_SLAVE
Operation not allowed on slave volume that is not consistent.
- SNAPSHOT_IS_FORMATTED
Snapshot is formatted
- NOT_ENOUGH_SPACE
No space to allocate for volume's current usage
- 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_INCOMPATIBLE_SIZE
Volume has different size than its snapshot
- VOLUME_TARGET_INVALID
The target volume is on the same tree as the base and source volumes, or the target volume is not of the same dynasty as the source and base volumes
- VOLUME_IS_COMPRESSED
Operation not allowed on compressed volume.

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 ]
[ compressed=<default|yes|no> ]
```

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

Name	Type	Description	Mandatory	Default
compressed	Enumeration	Must be specified as yes if the new volume should be compressed. Default value uses pool setting	N	default
ext_id	String	External identifier of the volume.	N	N/A

The name of the volume must be unique in the system.

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 soft capacity of the storage pool for the volume.

The command fails if the reservation cannot be committed.

The volume is logically formatted at creation time, which means that any read operation results in returning all zeros as a response.

The size is the actual "net" storage space, as seen by the user's applications, not including any mirroring or other data protection overhead.

The free space consumed by the volume will be the smallest multiple of 17GB which is bigger than the specified size.

The size can be specified either in gigabytes or in blocks (where each block is 512 bytes). If the size is specified in blocks, volumes are created in the exact size specified. If the size is specified in gigabytes, the actual volume size is rounded up to the nearest 17GB multiple (making the actual size identical to the free space consumed by the volume, as described above). This rounding up prevents a situation where storage space is not fully utilized because of a gap between the free space used and the space available to the application. The size specified in blocks is exact.

The term GB (gigabytes) is defined in this context as 10^9 (and not as 2^{30} as in many other contexts).

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 executed 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_CANNOT_HAVE_ZERO_SIZE
Volume size cannot be zero
- POOL_DOES_NOT_EXIST
Storage Pool does not exist
- VOLUME_EXISTS
Volume name already exists
- VOLUME_BAD_PREFIX
Volume name has a reserved prefix
- NOT_ENOUGH_SPACE
No space to allocate volume
- MAX_VOLUMES_REACHED
Maximum number of volumes already defined
- ELECTRONIC_LICENSE_NOT_APPROVED
Operation blocked until Electronic license approval

Troubleshooting: Please retrieve Electronic license version and accept it

- VOLUME_SIZE_ABOVE_LIMIT
Volume size specified is above limit
- INVALID_SLICE_OFFSET
Slice number out of range
- MAX_COMPRESSED_VOLUMES_REACHED
Maximum number of compressed volumes 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.
- ENCRYPTION_IN_PROGRESS
System is in the process of changing encryption activation state
- DOMAIN_MAX_VOLUMES_REACHED
The domain exceeds the maximum allowed number of volumes.
- COMPRESSED_VOLUME_TOO_SMALL
Volume size is below the minimum size (51GB) allowed for compressed volume
- COMPRESSED_VOLUME_TOO_BIG
Volume size is above the maximum size allowed for compressed volume
- COMPRESSED_CAPACITY_LIMIT_REACHED
Total capacity of compressed volumes exceeds the maximum limit
- COMPRESSION_REQUIRES_THIN_PROVISIONED_POOL
Compressed volumes require thin provisioned pools
- COMPRESSION_IS_NOT_ACTIVE
Compression was not enabled.

Creating multiple volumes

Use the **vol_create_many** command to create new volumes.

```
vol_create_many vol=<vol1[,vol2]...>
< size=<GB1[,GB2]...> | size_blocks=<BLOCKS1[,BLOCKS2]...>
> pool=PoolName [ ext_id=<id1[,id2]...> ]
[ compressed=<<default|yes|no>[,<default|yes|no>]...> ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	Volume names.	Y	N/A
size	Positive integer	Volume sizes in GB.	N	N/A
size_blocks	Positive integer	Sizes in the number of blocks.	N	N/A
compressed	Enumeration	Create volumes as compressed or not.	N	Use pool setting.
pool	Object name	The name of the storage pool to which the volume belongs.	Y	N/A
ext_id	String	External identifiers of the volumes.	N	No external id

Create multiple volumes with an efficient slice distribution. Can be used to create up to 10 volumes.

Example:

```
vol_create_many vol=managed_vol_1,managed_vol_2 size=2000,1500 pool=DBPool ext_id=id1,id2
```

Output:

```
Command executed 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

- ELECTRONIC_LICENSE_NOT_APPROVED
Operation blocked until Electronic license approval
Troubleshooting: Please retrieve Electronic license version and accept it
- ENCRYPTION_IN_PROGRESS
System is in the process of changing encryption activation state
- POOL_DOES_NOT_EXIST
Storage Pool does not exist

- **COMPRESSION_IS_NOT_ACTIVE**
Compression was not enabled.

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](#).

The volume is removed from all LUN maps that contain its mapping

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 its snapshots?

Return codes

- **VOLUME_BAD_NAME**
Volume name does not exist
- **VOLUME_HAS_MIRROR**
Mirror is defined for this volume
- **SNAPSHOT_IS_PART_OF_SNAPSHOT_GROUP**
Snapshot is part of a Snapshot Group
- **SNAPSHOT_IS_INTERNAL**
Internal snapshots cannot be mapped, modified or deleted.
- **VOLUME_BELONGS_TO_CG**
Volume belongs to a Consistency Group
- **VOLUME_IS_MAPPED**
Volume that is mapped to a host cannot be deleted

- **VOLUME_IS_BOUND**
Volume is bound to a ALU

Troubleshooting: Unbound the volume from the ALU
- **VOLUME_HAS_TRANSFORM**
Operation is not supported for volumes which are compressed/decompressed
- **VOLUME_HAS_MAPPED_SNAPSHOT**
Volume which has a snapshot that is mapped to a host cannot be deleted
- **SNAPSHOT_HAS_ACTIVE_SYNC_JOB**
Snapshot is currently a target of an active sync job

Troubleshooting: Please wait for 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**
IBM Hyper-Scale Mobility relation is defined for this volume
- **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.
- **VOLUME_HAS_DATA_MIGRATION**
Data Migration is defined for this volume
- **COMPRESSION_RECOVERY_IN_PROGRESS**
Operation cannot be completed because another recovery process is in progress

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 executed 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_FORMAT_VOLUME**
Volume *Volume* may contain data. Formatting will cause data loss. Are you sure you want to format volume *Volume*?

Return codes

- **VOLUME_BAD_NAME**
Volume name does not exist
- **VOLUME_HAS_SNAPSHOTS**
Volume has snapshots
- **VOLUME_IS_SNAPSHOT**
Operation is not permitted on snapshots
- **VOLUME_LOCKED**
Volume is locked
- **VOLUME_HAS_MIRROR**
Mirror is defined for this volume
- **VOLUME_HAS_DATA_MIGRATION**
Data Migration is defined for this volume
- **VOLUME_HAS_TRANSFORM**
Operation is not supported for volumes which are compressed/decompressed
- **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.
- **COMPRESSION_VOLUME_IS_OFFLINE**
Operation cannot be completed because the requested compressed volume is offline

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

This command lists volumes according to:

- Volume name
- Pool
- Consistency Group

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
master_name	Master Name	3
cg_name	Consistency Group	4
pool_name	Pool	5
creator	Creator	6

Field ID	Field output	Default position
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
serial	Serial Number	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
used_capacity	Used Capacity (GB)	9
used_capacity_MiB	Used Capacity (MiB)	N/A
capacity_used_by_snapshots_MiB	Capacity Used by Snapshots (MiB)	N/A
short_lived_io	Short Live IO	N/A
enabled_VAAI_cmd	VAAI commands enabled bitmask	N/A
user_disabled_VAAI	VAAI disabled by user	N/A
snapshot_format	Snapshot Format	N/A
compressed	Compressed	7
ratio	Compression Ratio (%)	8
saving	Compression Saving (GB)	10
saving_MiB	Compression Saving (MiB)	N/A
ssd_caching	SSD Caching State	N/A
use_ssd_caching_default	Use SSD Caching Default State	N/A
unmap_support	Unmap Support	N/A
managed	Managed	N/A
enable_unmap	unmap enabled	N/A
user_disabled_unmap	unmap disabled by user	N/A
marked	Marked	N/A
online	Online	N/A
inflated	Inflated	N/A
metadata_mismatch	MetaData Mismatch	N/A

Example:

vol_list

Output:

Name	Size (GB)	Master Name	Pool	Creator	Used Capacity (GB)
DBLog	3006		MainPool	admin	0
Dev	2010		MainPool	admin	0
Marketing	1013		MainPool	admin	0
Dev.snapshot_00001	2010	Dev	MainPool	admin	
Dev.snapshot_00002	2010	Dev	MainPool	admin	
Dev.snapshot_00003	2010	Dev	MainPool	admin	

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Allowed
Read-only users	Allowed
Technicians	Disallowed

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.

Example:

```
vol_lock vol=DBVolume
```

Output:

```
Command executed 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
Volume name does not exist
- SNAPSHOT_IS_PART_OF_SNAPSHOT_GROUP
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.
- VOLUME_HAS_TRANSFORM
Operation is not supported for volumes which are compressed/decompressed
- COMPRESSION_VOLUME_IS_OFFLINE
Operation cannot be completed because the requested compressed volume is offline

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.

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
Volume name does not exist
- VOLUME_EXISTS
Volume name already exists
- SNAPSHOT_IS_INTERNAL
Internal snapshots cannot be mapped, modified or deleted.
- SNAPSHOT_IS_PART_OF_SNAPSHOT_GROUP
Snapshot is part of a Snapshot Group
- VOLUME_BAD_PREFIX
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.
- 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_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.
- VOLUME_HAS_TRANSFORM
Operation is not supported for volumes which are compressed/decompressed

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

Name	Type	Description	Mandatory	Default
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 specified as an integer multiple of 10^9 bytes, but the actual new size of the volume is rounded up to the nearest valid size, which is an integer multiple of 16×2^{30} bytes.

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.

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
```

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).

Output:

```
Command executed 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_ENLARGE_VOLUME`
Are you sure you want to increase volume size?
- `ARE_YOU_SURE_YOU_WANT_TO_REDUCE_VOLUME`
Decreasing volume size may cause data loss. Are you sure you want to proceed?
- `VOLUME_WILL_CROSS_1TB_SIZE`
Many operating systems do not support a resize operation across the 1TB boundary, are you sure?

Troubleshooting: Snapshot backup before resize is advised

Return codes

- `VOLUME_BAD_NAME`
Volume name does not exist
- `ILLEGAL_VOLUME_SIZE`
Illegal volume size
- `NOT_ENOUGH_SPACE`
No space to allocate volume
- `REMOTE_VOLUME_SIZE_ABOVE_LIMIT`
Volume size specified is above limit of remote machine
- `VOLUME_LOCKED`

Volume is locked

- VOLUME_HAS_DATA_MIGRATION
Data Migration is defined for this volume
- CAN_NOT_SHRINK_MAPPED_VOLUME
Mapped volume's size cannot be decreased
- CAN_NOT_SHRINK_VOLUME_WITH_SNAPSHOTS
Size of volume with snapshots cannot be decreased
- CAN_NOT_SHRINK_REMOTE_VOLUME_WITH_SNAPSHOTS
Remote volume has snapshots
- CAN_NOT_SHRINK_MAPPED_REMOTE_VOLUME
Remote volume is mapped
- VOLUME_IS_BOUND
Volume is bound to a ALU

Troubleshooting: Unbound the volume from the ALU

- REMOTE_VOLUME_HAS_DATA_MIGRATION
Data Migration is defined for slave volume
- VOLUME_CANNOT_HAVE_ZERO_SIZE
Volume size cannot be zero
- CAN_NOT_SHRINK_SNAPSHOTS
Size of snapshots cannot be decreased
- CAN_NOT_RESIZE_ASYNC_INTERVAL_VOLUMES
Size of volumes with asynchronous mirroring cannot be changed
- CAN_NOT_SHRINK_VOLUME
Size of volumes cannot be decreased without explicit request
- MIRROR_SIZE_MISMATCH
Slave volume and Master Volume sizes are different
- MIRROR_POSSIBLE_SIZE_MISMATCH
Slave volume and Master Volume sizes may be different. This problem occurs whenever the Master does not receive an acknowledgment from the Slave until the command timed out, or any other unexpected failure.
- VOLUME_SIZE_ABOVE_LIMIT
Volume size specified is above limit
- CAN_NOT_SHRINK_COMPRESSED_VOLUME
Compressed volume's size cannot be decreased
- VOLUME_HAS_TRANSFORM
Operation is not supported for volumes which are compressed/decompressed
- COMMAND_NOT_SUPPORTED_FOR_OLVM_VOLUMES
This command is not supported for IBM Hyper-Scale Mobility volumes.
- MIRROR_IS_NON_OPERATIONAL
Mirror is non-operational
- VOLUME_IS_SLAVE
Volume is defined as a slave volume
- MIRROR_RETRY_OPERATION
There is an operation in progress on this mirror , please try again your request in a few seconds

Troubleshooting: Please try again the command in a few seconds

- VOLUME_HAS_MULTIPLE_MIRRORS
Volume has multiple mirrors, operation not allowed or target must be specified
- REMOTE_MIRROR_IS_STANDBY
Remote mirror is marked as standby
- COMPRESSED_CAPACITY_LIMIT_REACHED
Total capacity of compressed volumes exceeds the maximum limit
- COMPRESSED_VOLUME_TOO_BIG
Volume size is above the maximum size allowed for compressed volume
- COMPRESSION_VOLUME_IS_OFFLINE
Operation cannot be completed because the requested compressed volume is offline

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 completion code is returned (see below).

Example:

```
vol_unlock vol=DBVolume
```

Output:

```
Command executed 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

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_UNLOCK_SNAPSHOT
Are you sure you want to unlock snapshot *Snapshot*?

Return codes

- VOLUME_BAD_NAME
Volume name does not exist
- VOLUME_IS_SLAVE
Volume is defined as a slave volume
- SNAPSHOT_IS_INTERNAL
Internal snapshots cannot be mapped, modified or deleted.
- SNAPSHOT_IS_PART_OF_SNAPSHOT_GROUP
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.
- COMPRESSION_VOLUME_IS_OFFLINE
Operation cannot be completed because the requested compressed volume is offline

Compressing volumes

Use the **vol_transform_create** command to compress a regular volume or decompress a compressed volume.

```
vol_transform_create vol=VolName mode=<compress|decompress> [ delete_source=<yes|no> ]  
[ force_on_marginal_ratio=<yes|no> ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	The volume to be transformed.	Y	N/A
mode	Enumeration	Direction of the compression: uncompressed to compressed or compressed to uncompressed.	Y	N/A
delete_source	Boolean	Defines whether to delete the source volume when the compression is completed successfully.	N	yes
force_on_marginal_ratio	Boolean	This parameter is required for a successful transformation of a volume if (1) the mode is <i>compress</i> , and (2) the expected compression ratio is either unknown or marginal.	N	no

This command is used to compress regular volumes or decompress compressed volumes. Only one transform can be performed at a time. Additional transforms are queued and scheduled for running on a first-in first-out basis.

The system emits the COMPRESSED_VOLUME_MARGINAL_COMPRESSION event if the transformation mode is *compress*, but the volume data is not suitable for compression. The event is issued in the following cases:

- When the volume's used capacity reaches 30%, and the compression ratio is 15% or smaller (that is, the ratio between written data and used capacity is 85% and above), the event is issued with the minor severity. This event is disabled by default.
- When the volume's used capacity reaches 50%, and the compression ratio is 15% or smaller, with hysteresis of 10%, the event is issued with the minor severity. This event is disabled by default.
- When the volume's used capacity reaches 70%, and the compression ratio is 15% or smaller, with hysteresis of 10%, the event is issued with the minor severity.
- When the volume's usage ratio reaches 80%, the event is issued with the major severity, regardless of the compression ratio.
- When the volume's usage ratio reaches 90%, the event is issued with the critical severity, regardless of the compression ratio.

If a COMPRESSED_VOLUME_MARGINAL_COMPRESSION event is emitted, consider decompressing or resizing the volume. If the volume is not resized or decompressed by the time COMPRESSED_VOLUME_MARGINAL_COMPRESSION is emitted with the critical severity, the volume might get locked.

Note that the COMPRESSED_VOLUME_MARGINAL_COMPRESSION event can be issued for any compressed volume, regardless of whether it was transformed from a non-compressed volume or created as compressed in the first place.

Example:

```
xcli vol_transform_create vol=myvolume mode=compress
```

Output:

```
Command executed 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_COMPRESS_THE_VOLUME
Are you sure you want to compress the volume?
- ARE_YOU_SURE_YOU_WANT_TO_DECOMPRESS_THE_VOLUME
Are you sure you want to decompress the volume?

Return codes

- VOLUME_BAD_NAME
Volume name does not exist
- VOLUME_IS_IN_REQUESTED_COMPRESSION_FORM
Compress transformation was requested while volume is already compressed, or decompress was requested while the volume is already uncompressed
- VOLUME_ALREADY_IN_TRANSFORM_QUEUE
Transformation request for selected volume already exists
- NOT_ENOUGH_SPACE
No space to allocate volume

- **VOLUME_HAS_MIRROR**
Mirror is defined for this volume
- **VOLUME_HAS_DATA_MIGRATION**
Data Migration is defined for this volume
- **VOLUME_IS_SNAPSHOT**
Operation is not permitted on snapshots
- **COMPRESSION_REQUIRES_THIN_PROVISIONED_POOL**
Compressed volumes require thin provisioned pools
- **COMPRESSED_CAPACITY_LIMIT_REACHED**
Total capacity of compressed volumes exceeds the maximum limit
- **DOMAIN_MAX_VOLUMES_REACHED**
The domain exceeds the maximum allowed number of volumes.
- **MAX_VOLUMES_REACHED**
Maximum number of volumes already defined
- **MAX_COMPRESSED_VOLUMES_REACHED**
Maximum number of compressed volumes already defined
- **COMPRESSION_IS_NOT_ACTIVE**
compression was not enabled.
- **VOLUME_HAS_OLVM**
IBM Hyper-Scale Mobility relation is defined for this volume
- **MAXIMUM_VOLUME_TRANSFORMS_REACHED**
Maximum volume transforms reached.
- **COMPRESSED_VOLUME_TOO_SMALL**
Volume size is below the minimum size (51GB) allowed for compressed volume
- **COMPRESSED_VOLUME_TOO_BIG**
Volume size is above the maximum size allowed for compressed volume
- **VOLUME_HAS_SNAPSHOTS**
Operation is not permitted on volume with snapshots
- **VOLUME_IS_BOUND**
Volume is bound to a ALU

Troubleshooting: Unbound the volume from the ALU
- **COMPRESSION_VOLUME_IS_OFFLINE**
Operation cannot be completed because the requested compressed volume is offline
- **EXPECTED_COMPRESSION_RATIO_IS_UNKNOWN_OR_MARGINAL**
Operation is not allowed as it will result in a compressed volume with unknown or marginal compression ratio.

Troubleshooting: Add the volume to the comprestimation queue or use the `force_on_marginal_ratio=yes` parameter.

Managing volume compression and decompression

Use the **vol_transform_update** command to control volume compression and decompression transformation requests.

```
vol_transform_update vol=VolName action=<suspend|resume|retry>
```

Parameters

Name	Type	Description	Mandatory
vol	Object name	The volume to be transformed.	Y
action	Enumeration	Transformation action. Allows suspending, resuming, canceling or retrying the transformation.	Y

Example:

```
xcli vol_transform_update vol=DBVolume action=suspend
```

Output:

```
Command executed 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

- VOLUME_BAD_NAME
Volume name does not exist
- VOLUME_IS_NOT_IN_TRANSFORM_QUEUE
Transformation request for selected volume does not exist
- ACTION_NOT_ALLOWED_IN_CURRENT_TRANSFORM_STATE
Action is not allowed in current transformation state
- COMPRESSION_VOLUME_IS_OFFLINE
Operation cannot be completed because the requested compressed volume is offline

Deleting a compression or decompression request

Use the **vol_transform_delete** command to delete a volume compression or decompression request from the queue.

```
vol_transform_delete vol=VolName
```

Parameters

Name	Type	Description	Mandatory
vol	Object name	Volume which is being transformed.	Y

If the request is already running, this command will stop the process.

Example:

```
xcli vol_transform_delete vol=DBVolume
```

Output:

```
Command executed 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_ABORT_VOLUME_COMPRESSION
Are you sure you want to abort volume compression?
- ARE_YOU_SURE_YOU_WANT_TO_ABORT_VOLUME_DECOMPRESSION
Are you sure you want to abort volume decompression?

Return codes

- VOLUME_BAD_NAME
Volume name does not exist
- VOLUME_IS_NOT_IN_TRANSFORM_QUEUE
Transformation request for selected volume does not exist

Listing the volume compression activity

Use the **vol_transform_list** command to list the state of all volume transformation requests.

```
vol_transform_list [ vol=VolName ] [ domain=DomainName ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	The name of a specific volume for which a transformation request was initiated.	N	All volumes.
domain	Object name	The domain name.	N	All Domains

The possible states are: Pending, Running, Completed, Suspended, and Failed.

Field ID	Field output	Default position
name	Name	1
mode	Mode	2
status	Status	3
status_change_time	In current status since	4
progress	Progress (% completed)	5
estimated_complete_duration	Estimated time to completion (seconds)	6
size	Size (GB)	7
capacity	Used Capacity (GB)	8
delete_source	Delete Source?	9
position	Position	10

Example:

```
xcli vol_transform_list
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Allowed
Read-only users	Allowed
Technicians	Disallowed

Setting a volume's external identifier

Use the **vol_set_external_id** command to set a volume's external identifier.

```
vol_set_external_id vol=VolName ext_id=Identifier
```

Parameters

Name	Type	Description	Mandatory
vol	Object name	Volume name.	Y
ext_id	String	External identifier of the volume.	Y

Example:

```
vol_set_external_id vol=vol1  
ext_id=" _home_markdown_jenkins_workspace_Transform_in_STJTAG_Gen3_XCLI_r_command_volume_vol_set_external_id_myId"
```

Output:

```
Command executed 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

- VOLUME_BAD_NAME
Volume name does not exist

Clearing a volume's external identifier

Use the **vol_clear_external_id** command to clear a volume's external identifier.

```
vol_clear_external_id vol=VolName
```

Parameters

Name	Type	Description	Mandatory
vol	Object name	Volume name.	Y

Example:

```
vol_clear_external_id vol=vol1
```

Output:

```
Command executed 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

- VOLUME_BAD_NAME
Volume name does not exist
- VOLUME_IS_BOUND
Volume is bound to a ALU

Troubleshooting: Unbound the volume from the ALU

Displaying the allocated partitions of a volume

Use the **vol_allocation_bitmap** command to display the allocated partitions of a volume.

```
vol_allocation_bitmap vol=VolName offset=Offset count=N
```

Parameters

Name	Type	Description	Mandatory
vol	Object name	Name of the volume.	Y
count	Integer	Number of partitions to check.	Y
offset	Integer	First partition to check.	Y

Displays the allocated partitions of a volume starting from the specified partition. The allocated count partitions are shown as an array of integers, each integer representing a bitmap of 8 partitions.

Example:

```
vol_allocation_bitmap vol=myVol offset=10 count=120
```

Output:

```
TBD
```

Access control

User Category	Permission

User Category	Permission
Storage administrator	Disallowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- **VOLUME_BAD_NAME**
Volume name does not exist
- **INVALID_COUNT**
The offset plus count is beyond the volume address space
- **INVALID_OFFSET**
Specified partition offset is beyond the volume address space.
- **VOLUMES_TO_COMPARE_ARE_NOT_RELATED**
The source volume and target volume are not of the same dynasty and therefore XIV cannot compute the difference between them.
- **VOLUME_IS_COMPRESSED**
Operation not allowed on compressed volume.

Displaying the difference of allocated partitions between two volumes

Use the **vol_diff** command to display the allocated partitions' difference between two volumes.

```
vol_diff vol_src=VolName vol_trg=VolName offset=Offset count=N
```

Parameters

Name	Type	Description	Mandatory
vol_src	Object name	Name of the first volume.	Y
vol_trg	Object name	Name of the second volume.	Y
count	Integer	Number of partitions to check.	Y
offset	Integer	First partition to check.	Y

The result shows the different count partitions from the starting partition as an array of integers, each integer representing a bitmap of 8 partitions.

Example:

```
vol_diff vol_src=v1 vol_trg=v2 offset=10 count=120
```

Output:

TBD

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

- **SOURCE_VOLUME_BAD_NAME**
Source volume name does not exist
- **TARGET_VOLUME_BAD_NAME**
Target volume name does not exist
- **VOLUMES_TO_COMPARE_ARE_NOT_RELATED**
The source volume and target volume are not of the same dynasty and therefore XIV cannot compute the difference between them.
- **INVALID_COUNT**

The offset plus count is beyond the volume address space

- **INVALID_OFFSET**
Specified partition offset is beyond the volume address space.
- **VOLUME_IDENTICAL**
Source and target are the same volume
- **VOLUME_IS_COMPRESSED**
Operation not allowed on compressed volume.

Triggering compression estimation for a system volume

Use the **vol_comprestimate** command to request an estimate of a volume's compression ratio.

```
vol_comprestimate vol=VolName
```

Parameters

Name	Type	Description	Mandatory
vol	Object name	Name of the volume on which to trigger the compression estimation.	Y

If continuous estimation is activated, this command moves the specified volume to the head of the queue, before the automatically scheduled volumes and after the last pending user-specified volume.

Example:

```
vol_comprestimate vol=MyVol
```

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

- **COMPRESSED_VOLUME_TOO_SMALL**
Volume size is below the minimum size (51GB) allowed for compressed volume
- **VOLUME_IS_SNAPSHOT**
Operation is not permitted on snapshots
- **VOLUME_BAD_NAME**
Volume name does not exist
- **COMPRESSED_VOLUME_TOO_BIG**
Volume size is above the maximum size allowed for compressed volume
- **VOLUME_IS_COMPRESSED**
Operation not allowed on compressed volume.
- **VOLUME_ALREADY_IN_QUEUE_FOR_COMPRESSION_ESTIMATION**
Volume is already in queue for compression estimation.
- **VOLUME_HAS_TRANSFORM**
Operation is not supported for volumes which are compressed/decompressed

Removing a volume from the compression estimation queue

Use the **vol_comprestimate_stop** command to remove the specified volume from the compression estimation queue.

```
vol_comprestimate_stop vol=VolName
```

Parameters

Name	Type	Description	Mandatory
vol	Object name	Volume name	Y

Once removed from the compression estimation queue, the volume will not be processed by the compression estimator unless it is re-inserted into the queue via the **vol_comprestimate** command.

Example:

```
vol_comprestimate_stop vol=MyVol
```

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

- COMPRESSED_VOLUME_TOO_SMALL
Volume size is below the minimum size (51GB) allowed for compressed volume
- VOLUME_IS_SNAPSHOT
Operation is not permitted on snapshots
- VOLUME_BAD_NAME
Volume name does not exist
- COMPRESSED_VOLUME_TOO_BIG
Volume size is above the maximum size allowed for compressed volume
- VOLUME_IS_COMPRESSED
Operation not allowed on compressed volume.

Displaying the volume's estimated compression ratio

Use the **vol_comprestimate_list** to display the volume's compression estimates.

```
vol_comprestimate_list [ vol=VolName ] [ domain=DomainName ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	Name of a specific volume to be listed.	N	All volumes
domain	Object name	The domain name.	N	All Domains

This command displays the estimated compression ratio for either a specified volume or for all volumes.

Example:

```
vol_comprestimate_list
```

Output:

Name	Compression Ratio (%)	Last estimation time	Index	Status
v1			1	In Progress
v2			2	Pending

Field ID	Field output	Default position
name	Name	1
ratio	Compression Ratio (%)	2
last_estimation_time	Last estimation time	3
status	Status	4

Field ID	Field output	Default position
position	Position	5

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Allowed
Read-only users	Allowed
Technicians	Disallowed

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.
- [Creating multiple snapshots](#)
Use the **snapshot_create_many** command to create multiple snapshots 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**
Operation is permitted only on snapshots
- **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**
Snapshot is part of a Snapshot Group
- **VOLUME_BAD_NAME**
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.

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 volume 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 executed 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

- VOLUME_BAD_NAME
Volume name does not exist
- MAX_VOLUMES_REACHED
Maximum number of volumes already defined
- 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
Operation is not permitted on snapshots
- VOLUME_EXISTS
Volume name already exists
- VOLUME_BAD_PREFIX
Volume name has a reserved prefix
- VOLUME_HAS_TRANSFORM
Operation is not supported for volumes which are compressed/decompressed
- VOLUME_DATA_MIGRATION_UNSYNCHRONIZED
Data Migration has not completed to this volume
- OVERWRITE_SNAPSHOT_BAD_NAME
Snapshot name does not exist
- OVERWRITE_SNAPSHOT_IS_MASTER_VOL
This snapshot cannot be overwritten as it is a master volume.

- **SNAPSHOT_OVERWRITE_MISMATCH**
Specified snapshot is not a snapshot of the specified volume
 - **SNAPSHOT_IS_PART_OF_SNAPSHOT_GROUP**
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**
Operation not allowed on slave volume that is not consistent.
 - **SNAPSHOT_HAS_ACTIVE_SYNC_JOB**
Snapshot is currently a target of an active sync job
- Troubleshooting:** Please wait for sync job to complete
- **TOO_MANY_FAST_SNAPSHOTS_IN_VOLUME**
Max number of fast snapshots for this volume already 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.
 - **NUM_VOLUMES_WILL_EXCEED_MAXIMUM**
Can not create all the volumes, as the num of volumes will exceed its maximum value.
 - **DOMAIN_WILL_EXCEED_MAXIMUM_VOLUMES_ALLOWED**
Can not create all of the volumes, as the domain will exceed the maximum allowed number of volumes.
 - **MAX_COMPRESSED_VOLUMES_REACHED**
Maximum number of compressed volumes already defined

Creating multiple snapshots

Use the **snapshot_create_many** command to create multiple snapshots of an existing volume.

```
snapshot_create_many vol=<vol1[,vol2]...> [ name=<name1[,name2]...> ]
[ delete_priority=del_value ] [ ext_id=<id1[,id2]...> ]
```

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
ext_id	String	External identifier of the volume.	N	No external id

This command creates a single snapshot per each of the listed volumes.

Example:

```
snapshot_create_many vol=v1,v2 name=snap1,snap2 ext_id=id1,id2
```

Output:

Command executed 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

- **ALL_VOLUMES_MUST_BE_FROM_SAME_DOMAIN**
All the volumes in a single command must belong to a single domain.
- **NUM_VOLUMES_WILL_EXCEED_MAXIMUM**

Can not create all the volumes, as the num of volumes will exceed its maximum value.

- DOMAIN_WILL_EXCEED_MAXIMUM_VOLUMES_ALLOWED
Can not create all of the volumes, as the domain will exceed the maximum allowed number of volumes.

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
Volume name does not exist
- VOLUME_IS_NOT_A_SNAPSHOT
Operation is permitted only on snapshots
- SNAPSHOT_IS_INTERNAL
Internal snapshots cannot be mapped, modified or deleted.
- SNAPSHOT_IS_PART_OF_SNAPSHOT_GROUP
Snapshot is part of a Snapshot Group
- SNAPSHOT_IS_MAPPED
Snapshot that is mapped to a host cannot be deleted
- VOLUME_IS_BOUND
Volume is bound to a ALU
Troubleshooting: Unbound the volume from the ALU
- SNAPSHOT_HAS_ACTIVE_SYNC_JOB
Snapshot is currently a target of an active sync job
Troubleshooting: Please wait for 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.

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 executed 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
Volume name does not exist
- MAX_VOLUMES_REACHED
Maximum number of volumes already defined
- MAX_COMPRESSED_VOLUMES_REACHED
Maximum number of compressed volumes already defined
- DOMAIN_MAX_VOLUMES_REACHED
The domain exceeds the maximum allowed number of volumes.
- VOLUME_IS_NOT_A_SNAPSHOT
Operation is permitted only on snapshots
- VOLUME_EXISTS
Volume name already exists
- SNAPSHOT_IS_PART_OF_SNAPSHOT_GROUP
Snapshot is part of a Snapshot Group
- VOLUME_BAD_PREFIX
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.

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 executed 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**
Snapshot is part of a Snapshot Group
- **VOLUME_BAD_NAME**
Volume name does not exist
- **SNAPSHOT_IS_INTERNAL**
Internal snapshots cannot be mapped, modified or deleted.
- **MAX_VOLUMES_REACHED**
Maximum number of volumes already defined
- **SNAPSHOT_IS_FORMATTED**
Snapshot is formatted
- **ELCS_CANNOT_BE_FORMATTED**
The snapshot is an ELCS and cannot be formatted.
- **VOLUME_IS_NOT_A_SNAPSHOT**
Operation is permitted only on snapshots
- **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 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
master_name	Master Name	3
cg_name	Consistency Group	4
pool_name	Pool	5
creator	Creator	6
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
serial	Serial Number	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
used_capacity	Used Capacity (GB)	9
used_capacity_MiB	Used Capacity (MiB)	N/A
capacity_used_by_snapshots_MiB	Capacity Used by Snapshots (MiB)	N/A
short_lived_io	Short Live IO	N/A
enabled_VAAI_cmd	VAAI commands enabled bitmask	N/A
user_disabled_VAAI	VAAI disabled by user	N/A
snapshot_format	Snapshot Format	N/A
compressed	Compressed	7
ratio	Compression Ratio (%)	8
saving	Compression Saving (GB)	10
saving_MiB	Compression Saving (MiB)	N/A
ssd_caching	SSD Caching State	N/A
use_ssd_caching_default	Use SSD Caching Default State	N/A
unmap_support	Unmap Support	N/A
managed	Managed	N/A
enable_unmap	unmap enabled	N/A
user_disabled_unmap	unmap disabled by user	N/A
marked	Marked	N/A
online	Online	N/A
inflated	Inflated	N/A
metadata_mismatch	MetaData Mismatch	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
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**
Volume name does not exist
- **VOLUME_IS_NOT_A_SNAPSHOT**
Operation is permitted only on snapshots
- **NOT_ENOUGH_SPACE**
No space to allocate volume
- **SNAPSHOT_IS_PART_OF_SNAPSHOT_GROUP**
Snapshot is part of a Snapshot Group
- **VOLUME_HAS_MIRROR**
Mirror is defined for this volume
- **VOLUME_LOCKED**
Volume is locked
- **SNAPSHOTS_BELONG_TO_DIFFERENT_MASTERS**
Target snapshot and source snapshot should be snapshots of the same volume
- **TARGET_SNAPSHOT_BAD_NAME**
Target snapshot name does not exist
- **TARGET_SNAPSHOT_IS_PART_OF_SNAPSHOT_GROUP**
Target snapshot is part of a Snapshot Group
- **TARGET_SNAPSHOT_IS_MASTER**
Target snapshot is a master volume
- **TARGET_SNAPSHOT_SAME_AS_SNAPSHOT**
Source snapshot cannot be the target snapshot
- **VOLUME_HAS_TRANSFORM**
Operation is not supported for volumes which are compressed/decompressed
- **TARGET_SNAPSHOT_HAS_ACTIVE_SYNC_JOB**
Target snapshot is currently a 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.
- **COMPRESSION_VOLUME_IS_OFFLINE**
Operation cannot be completed because the requested compressed volume is offline

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.
- [Suspending I/O execution on a consistency group](#)
Use the **io_pause** command to suspend I/O execution on a consistency group.
- [Resuming I/O execution](#)
Use the **io_resume** command to resume I/O execution on a consistency group, previously suspended with the **io_pause** command.
- [Listing the status of consistency groups with paused I/O](#)
Use the **io_pause_list** command to list the status of consistency groups for which the **io_pause** command was invoked.
- [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 a consistency group from a cross-system consistency group](#)
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 from a cross-system consistency group](#)
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

This command adds a volume to a consistency group. The consistency group can contain up to 128 volumes.

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.

In case of acknowledgment timeout:

- If the command is issued on a mirrored master consistency group, which does not receive an acknowledgment from the slave until the command times out or due to an unexpected failure, a completion code is returned (MIRROR_POSSIBLE_CONS_GROUP_MEMBERSHIP_MISMATCH).

If the command **cg_add_vol** is issued on a mirrored master consistency group, which does not receive an acknowledgment from the slave until the command times out or due to an unexpected failure, a new completion code is returned (MIRROR_POSSIBLE_CONS_GROUP_MEMBERSHIP_MISMATCH, meaning that the member lists of the mirror consistency group peers might not be the same).

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**
Volume Mirror has different designation than Consistency Group Mirror. Are you sure you want to add that Volume to that CG?

Return codes

- **VOLUME_BAD_NAME**
Volume name does not exist
- **CONS_GROUP_BAD_NAME**
Consistency Group name does not exist.
- **CONS_GROUP_IS_SLAVE**
Consistency Group is mirroring slave.
- **MAX_VOLUMES_IN_CONS_GROUP_REACHED**
The maximum permitted number of volumes per Consistency Group is reached.

Consistency Group contains maximum number of volumes.
- **MAX_VOLUMES_IN_REMOTE_CONS_GROUP_REACHED**
The maximum permitted number of volumes per the remote Consistency Group is reached.

Remote Consistency Group contains maximum number of volumes.
- **MIRROR_HAS_SYNC_JOB**
Operation is not permitted on a mirror with active sync jobs
- **MIRROR_IS_NOT_SYNCHRONIZED**
Mirror is not synchronized
- **MIRROR_LAST_SYNC_TIMES_DIFFER**
All mirrors should have the same last sync time.
- **MIRROR_RETRY_OPERATION**
There is an operation in progress on this mirror , please try again your request in a few seconds

Troubleshooting: Please try again the command in a few seconds
- **REMOTE_VOLUME_BAD_POOL**
Remote volume and remote Consistency Group belong to different Storage Pools
- **REMOTE_VOLUME_BELONGS_TO_CONS_GROUP**
Remote Volume belongs to a Consistency Group
- **TARGET_NOT_CONNECTED**
There is currently no connection to the target system
- **VOLUME_BAD_POOL**
Volume belongs to a different Storage Pool
- **VOLUME_BELONGS_TO_CG**
Volume belongs to a Consistency Group
- **VOLUME_DATA_MIGRATION_UNSYNCHRONIZED**
Data Migration has not completed to this volume
- **VOLUME_IS_SNAPSHOT**
Operation is not permitted on snapshots
- **CONS_GROUP_MIRROR_SCHEDULE_MISMATCH**
Volumes under Consistency Group Mirror should have the same mirroring schedule.
- **CONS_GROUP_MIRROR_TARGET_MISMATCH**
Volumes under Consistency Group Mirror should have the same mirroring Target.
- **CONS_GROUP_MIRROR_ROLE_MISMATCH**
Volumes under Consistency Group Mirror should have the same mirroring role.

- **CONS_GROUP_MIRROR_ACTIVATION_MISMATCH**
Volumes under Consistency Group Mirror should have the same mirroring activation state.
- **REMOTE_CONS_GROUP_MIRROR_SCHEDULE_MISMATCH**
Volumes under Consistency Group Mirror in remote machine should have identical mirroring schedule.
- **CONS_GROUP_MIRROR_TYPE_MISMATCH**
Volumes under Consistency Group Mirror should have the same mirroring type.
- **MIRROR_POSSIBLE_CONS_GROUP_MEMBERSHIP_MISMATCH**
Mirrored CG may contain different volumes on Master and Slave. This problem occurs whenever the `cg_add_vol` command results with the Master not receiving an acknowledgment from the Slave until the command timed out, or any other unexpected failure.
- **REMOTE_CONS_GROUP_APPLICATION_CONSISTENCY_MISMATCH**
Application consistency of the volume doesn't match the state of other volumes in the group on the remote machine.
- **CONS_GROUP_APPLICATION_CONSISTENCY_MISMATCH**
Application consistency of the volume doesn't match the state of other volumes in the group.
- **VOLUME_HAS_TRANSFORM**
Operation is not supported for volumes which are compressed/decompressed
- **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**
IBM Hyper-Scale Mobility relation is defined for this volume
- **VOLUME_HAS_MULTIPLE_MIRRORS**
volume has multiple mirrors, operation not allowed or target must be specified
- **REMOTE_MIRROR_IS_STANDBY**
remote mirror is marked as standby

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
gp	Object name	Grouped Pool of the Consistency Group.	N

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 system can contain up to 256 consistency groups.

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.

Example:

```
cg_create pool=p_1 cg=DBgroup
```

Output:

```
Command executed 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**
Consistency Group name already exists.
- **MAX_CONS_GROUPS_REACHED**
Maximum number of Consistency Groups already defined.
- **POOL_DOES_NOT_EXIST**
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.
- **GROUPED_POOL_DOES_NOT_EXIST**
Grouped Pool does not exist.

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 is 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**
Consistency Group name does not exist.
- **CONS_GROUP_NOT_EMPTY**
This operation is only allowed on an empty Consistency Group.
- **CONS_GROUP_HAS_MIRROR**
Consistency Group has mirroring defined for it.
- **CONS_GROUP_BELONGS_TO_XCG**
Consistency Group belongs to another Cross 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.

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.

Field ID	Field output	Default position
name	Name	1
pool	Pool Name	2
gp_based	GP Based	N/A
mirrored	Mirrored	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 consistency 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).

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
Volume name does not exist
- VOLUME_NOT_IN_CONS_GROUP
Volume does not belong to a Consistency Group
- TARGET_NOT_CONNECTED
There is currently no connection to the target system
- VOLUME_IS_SNAPSHOT
Operation is not permitted on snapshots
- CONS_GROUP_IS_SLAVE
Consistency Group is mirroring slave.
- MIRROR_RETRY_OPERATION
There is an operation in progress on this mirror , please try again your request in a few seconds
Troubleshooting: Please try again the command in a few seconds
- MIRROR_HAS_SYNC_JOB
Operation is not permitted on a mirror with active sync jobs
- MIRROR_POSSIBLE_CONS_GROUP_MEMBERSHIP_MISMATCH
Mirrored CG may contain different volumes on Master and Slave. This problem occurs whenever the cg_add_vol command results with the Master not receiving an acknowledgment from the Slave until the command timed out, or any other unexpected failure.
- VOLUME_IS_NOT_CONSISTENT_SLAVE
Operation not allowed on slave volume that is not consistent.
- SNAPSHOT_HAS_ACTIVE_SYNC_JOB
Snapshot is currently a target of an active sync job
Troubleshooting: Please wait for sync job to complete
- VOLUME_HAS_TRANSFORM
Operation is not supported for volumes which are compressed/decompressed
- 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
IBM Hyper-Scale Mobility relation is defined for this volume
- REMOTE_MIRROR_IS_STANDBY
Remote mirror is marked as standby

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
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- **CONS_GROUP_BAD_NAME**
Consistency Group name does not exist.
- **CONS_GROUP_NAME_EXISTS**
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.

Suspending I/O execution on a consistency group

Use the **io_pause** command to suspend I/O execution on a consistency group.

```
io_pause cg=cgName [ milli_seconds_to_resume=MilliSecondsTimeout ] [ allow_read=AllowRead ]
```

Parameters

Name	Type	Description	Mandatory	Default
milli_seconds_to_resume	Positive integer	Timeout for auto resume. The measurement starts when current I/O execution on the consistency group is completed.	N	10000
allow_read	Boolean	Controls whether to continue reading I/Os while I/Os are suspended.	N	yes
cg	Object name	CG name	Y	N/A

The I/O execution is suspended with an auto-resume timeout.

It is possible to suspend execution of all I/Os or writes only.

Example:

```
io_pause cg=test_cg milli_seconds_to_resume=10000
```

Output:

```
command:
  code = "SUCCESS"
  status = "0"
  status_str = "Command completed successfully"
```

```
return:
    token_id = "6343971831808"
```

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

- **COMMAND_IS_NOT_VALID_IN_CURRENT_SYSTEM_STATE**
The requested command cannot be invoked in the current system state
- **PAUSE_IO_TIMEOUT_OUT_OF_RANGE**
Timeout parameter is out of range
- **TOO_MANY_IO_PAUSE_ISSUED**
Too many Pause IOs are in progress
- **CONS_GROUP_BAD_NAME**
Consistency Group name does not exist.
- **CONS_GROUP_IS_SLAVE**
Consistency Group is mirroring slave.
- **IO_PAUSE_ALREADY_ISSUED_FOR_CONS_GROUP**
Volume(s) belonging to the Consistency Group are already paused
- **VOLUME_HAS_TRANSFORM**
Operation is not supported for volumes which are compressed/decompressed

Resuming I/O execution

Use the **io_resume** command to resume I/O execution on a consistency group, previously suspended with the **io_pause** command.

```
io_resume token_id=Token
```

Parameters

Name	Type	Description	Mandatory
token_id	Positive integer	The token returned by the io_resume command.	Y

Example:

```
io_resume token_id=6343971831808
```

Output:

```
command:
    code = "SUCCESS"
    status = "0"
    status_str = "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.

User Category	Permission	Condition
Security administrator	Disallowed	N/A
Read-only users	Disallowed	N/A
Technicians	Disallowed	N/A

Return codes

- **CONS_GROUP_IS_NOT_PAUSED**
Consistency Group is not paused or auto-resume timeout expired
- **CONS_GROUP_DEFINITION_MODIFIED_DURING_IO_PAUSE**
Consistency Group definitions changed during pause io period

Listing the status of consistency groups with paused I/O

Use the **io_pause_list** command to list the status of consistency groups for which the **io_pause** command was invoked.

```
io_pause_list [ token_id=Token ]
```

Parameters

Name	Description	Mandatory	Default
token_id	Optional filter value to show the status for a specific token. 0 means that the filter is not applied.	N	0

This command displays the detailed status of the consistency groups on which I/O was suspended with the **io_pause** command.

Example:

```
io_pause_list
```

Output:

```
command:
  code = "SUCCESS"
  status = "0"
  status_str = "Command completed successfully"
  return:
    stop_io 0:
      allow_read = "yes"
      cg_name = "cg_test"
      config_changed = "no"
      inode_list_changed = "no"
      num_volumes = "1"
      resume_pending = "no"
      stop_io_elapsed_time = "4062"
      timeout = "10000"
      token = "6343971831808"
```

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

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
Cross Consistency Group name already exists.
- MAX_XCGS_REACHED
Maximum number of Cross Consistency Groups already defined.

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
Cross Consistency Group name does not exist.
- MAX_CONS_GROUPS_IN_XCG_REACHED
Cross Consistency Group contains maximum number of cgs.
- CONS_GROUP_IS_SLAVE
Consistency Group is mirroring slave.
- CONS_GROUP_BAD_NAME
Consistency Group name does not exist.
- CONS_GROUP_ALREADY_IN_XCG
Consistency Group already belongs to Cross Consistency Group.
- CONS_GROUP_BELONGS_TO_XCG
Consistency Group belongs to another Cross Consistency Group.

Removing a consistency group from a cross-system consistency group

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 cons group 'CG' from its cross Consistency Group?

Return codes

- XCG_BAD_NAME
Cross Consistency Group name does not exist.
- CONS_GROUP_BAD_NAME
Consistency Group name does not exist.
- XCG_IS_EMPTY
Consistency Group is empty.
- CONS_GROUP_NOT_IN_XCG
Consistency Group doesnt belong to Cross 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
Cross Consistency Group name does not exist.
- MAX_REMOTE_SYSTEMS_IN_XCG_REACHED
Cross Consistency Group contains maximum number of remote systems.
- REMOTE_SYSTEM_ALREADY_ADDED
Remote system belongs to Cross Consistency Group

Removing a remote system from a cross-system consistency group

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

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	Disallowed	N/A
Technicians	Disallowed	N/A

Return codes

- XCG_BAD_NAME
Cross Consistency Group name does not exist.
- REMOTE_SYSTEM_NOT_IN_XCG
Remote system doesn't belong to Cross 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
Cross 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=XcGroup1
```

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
Cross 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
Cross Consistency Group name does not exist.
- XCG_NOT_EMPTY
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
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

Grouped pool management commands

This chapter describes the grouped pool management commands.

- [Creating grouped pools](#)
Use the **gp_create** command to create a grouped pool.
- [Listing grouped pools](#)
Use the **gp_list** command to list grouped pools.
- [Deleting a grouped pool](#)
Use the **gp_delete** command to delete a grouped pool.
- [Renaming grouped pools](#)
Use the **gp_rename** command to rename a grouped pool.
- [Moving capacity in a grouped pool](#)
Use the **gp_capacity_shift** command to move capacity between pools in a grouped pool.

Creating grouped pools

Use the **gp_create** command to create a grouped pool.

```
gp_create gp=gpName thin_pool=thinPoolName thick_pool=thickPoolName  
[ meta_pool=metaPoolName ]
```

Parameters

Name	Type	Description	Mandatory	Default
gp	Object name	Name of the grouped pool.	Y	N/A
thin_pool	Object name	Thin pool of the grouped pool.	Y	N/A
thick_pool	Object name	Thick pool of the grouped pool.	Y	N/A
meta_pool	Object name	Meta pool of the grouped pool.	N	None

This command creates a grouped pool. Grouped pools enable consistency group consisting of both thin and thick volumes.

The name of the grouped pool must be unique in the system. The system is capable of containing up to 128 grouped pools.

The thick and thin pools of the grouped pool must be specified.

Example:

```
gp_create meta_pool=p_1 thin_pool=p_2 thick_pool=p_3 gp=gp_1
```

Output:

```
Command executed 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

- GROUPED_POOL_NAME_ALREADY_EXISTS
Grouped Pool name already exists.
- MAX_GROUPED_POOLS_REACHED
Maximum number of Grouped Pools already defined.
- POOL_DOES_NOT_EXIST
Storage Pool does not exist
- GROUPED_POOL_THIN_MUST_BE_THIN
Thin pool must be thin provisioned.
- THIN_POOL_ALREADY_BELONGS_TO_A_GROUPED_POOL
Thin pool belongs to a Grouped Pool.
- THICK_POOL_ALREADY_BELONGS_TO_A_GROUPED_POOL
Thick pool belongs to a Grouped Pool.
- META_POOL_ALREADY_BELONGS_TO_A_GROUPED_POOL
Metadata pool belongs to a Grouped Pool.
- ALL_POOLS_MUST_BE_UNIQUE
The Meta Pool, Thick and Thin pools must be unique
- GROUPED_POOL_NOT_ALL_POOLS_BELONG_TO_SAME_DOMAIN
Not all pools belongs to the same domain.
- DOMAIN_MAX_GROUPED_POOLS_REACHED
Domain exceed the max number of grouped pools allowed.
- NO_GROUPED_POOL_ALLOWED_IN_DOMAIN
The domain definition does not allow grouped pools.

Listing grouped pools

Use the **gp_list** command to list grouped pools.

```
gp_list [ gp=gpName ] [ domain=DomainName ]
```

Parameters:

Name	Type	Description	Mandatory	Default
gp	Object name	The name of a consistency group.	N	All
domain	Object name	The domain name.	N	All Domains

This command lists select details for all grouped pools. If a grouped pool name is specified, it is the only one to be listed.

Listed details include the following:

- Name
- Meta pool name
- Thin pool name
- Thick pool name

Field ID	Field output	Default position
name	Name	1
meta_pool	Meta pool name	2
thin_pool	Thin pool name	3
thick_pool	Thick pool name	4

Example:

```
gp_list gp=gp_1
```

Output:

Name	Meta Pool Name	Thin Pool Name	Thick Pool Name
gp_1	metaPool	thinPool	thickPool

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Allowed
Read-only users	Allowed
Technicians	Allowed

Deleting a grouped pool

Use the **gp_delete** command to delete a grouped pool.

```
gp_delete gp=gpName
```

Parameters

Name	Type	Description	Mandatory
gp	Object name	Name of the grouped pool to be deleted.	Y

Example:

```
gp_delete gp=gp_1
```

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

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_OPERATE_ON_MANAGED_OBJECT

This is a managed object. Performing manual operations on it may cause severe problems to the managing software. Are you sure you want to perform the operation on this managed object?

Return codes

- GROUPED_POOL_DOES_NOT_EXIST
Grouped Pool does not exist.
- GROUPED_POOL_HAS_CG
Grouped Pool has Consistency Groups defined
- GROUPED_POOL_THICK_POOL_HAS_GOLDEN_SNAPSHOTS
The thick pool of the Grouped Pool has golden snapshots.
- 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.

Renaming grouped pools

Use the **gp_rename** command to rename a grouped pool.

```
gp_rename gp=gpName new_name=Name
```

Parameters

Name	Type	Description	Mandatory
gp	Object name	The name of the grouped pool to be renamed.	Y
new_name	Object name	The new name of the grouped pool.	Y

The new name of the group pool must be unique in the system.

This command succeeds even if the new name is identical to the current name.

Example:

```
gp_rename gp=gp_1 new_name=gp_2
```

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

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_OPERATE_ON_MANAGED_OBJECT
This is a managed object. Performing manual operations on it may cause severe problems to the managing software. Are you sure you want to perform the operation on this managed object?

Return codes

- GROUPED_POOL_DOES_NOT_EXIST
Grouped Pool does not exist.
- GROUPED_POOL_NAME_ALREADY_EXISTS
Grouped Pool 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.

Moving capacity in a grouped pool

Use the **gp_capacity_shift** command to move capacity between pools in a grouped pool.

```
gp_capacity_shift gp=gpName src_role=<thin|thick|meta>
dst_role=<thin|thick|meta> capacity=GB
```

Parameters

Name	Type	Description	Mandatory
gp	Object name	The name of the grouped pool.	Y
src_role	Positive integer	The pool type from which the capacity should be removed.	Y
dst_role	Positive integer	The pool type to which the capacity should be added.	Y
capacity	Positive integer	Capacity size (in gigabytes) to be moved.	Y

Example:

```
gp_capacity_shift gp=gp_1 src_role=thin dst_role=thick capacity=200
```

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

Warnings

- ARE_YOU_SURE_YOU_WANT_TO_OPERATE_ON_MANAGED_OBJECT
This is a managed object. Performing manual operations on it may cause severe problems to the managing software. Are you sure you want to perform the operation on this managed object?

Return codes

- GROUPED_POOL_DOES_NOT_EXIST
Grouped Pool does not exist.
- GROUPED_POOL_HAS_NO_META_POOL
Grouped Pool has no Meta Pool defined
- HARD_SIZE_SMALLER_THAN_SNAPSHOT_SIZE
Snapshot size must be equal or smaller than hard size
- POOL_TOO_SMALL
Storage Pool usage exceeds requested size
- POOL_SOFT_TOO_SMALL
Requested soft size is smaller than the sum of sizes of volumes in the Storage Pool
- SOURCE_POOL_HARD_TOO_SMALL
The source pool usage exceeds the requested size.
- REACHED_POOL_MAX_HARD_CAPACITY
Reached max pool hard capacity
- POOL_MUST_BE_THIN
This pool has a golden snapshot and therefore must be thin provisioned.
- 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.
- GROUPED_POOL_THIN_MUST_BE_THIN
Thin pool must be thin provisioned.

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
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.
- 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
Maximum number of volumes already defined
- MAX_COMPRESSED_VOLUMES_REACHED
Maximum number of compressed volumes already defined
- DOMAIN_MAX_VOLUMES_REACHED
The domain exceeds the maximum allowed number of volumes.
- CONS_GROUP_BAD_NAME
Consistency Group name does not exist.
- SNAPSHOT_GROUP_BAD_NAME
Snapshot Group name does not exist.
- SNAPSHOT_GROUP_BAD_PREFIX
Snapshot Group name has a reserved prefix.
- SNAPSHOT_GROUP_NAME_EXISTS
Snapshot Group name already exists.
- CONS_GROUP_EMPTY
Operation is not allowed on an empty Consistency Group.
- CONS_GROUP_MISMATCH
Snapshot Group does not match Consistency Group volumes.
- OVERWRITE_SNAPSHOT_GROUP_DOES_NOT_BELONG_TO_GIVEN_GROUP
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
Operation not allowed on slave volume that is not consistent.
- SNAPSHOT_GROUP_IS_INTERNAL

Internal snapshots cannot be mapped, modified in any way or deleted.

- **SNAPSHOT_GROUP_ILLEGAL_PRIORITY**
Illegal snapshot group priority; must be an integer between 1 and 4.
- **SNAPSHOT_HAS_ACTIVE_SYNC_JOB**
Snapshot is currently a 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.
- **CONS_GROUP_TOKEN_MISMATCH**
Token does not match Consistency Group.
- **VOLUME_HAS_TRANSFORM**
Operation is not supported for volumes which are compressed/decompressed

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
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**
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 in any way or deleted.

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
Snapshot Group name does not exist.
- SNAPSHOT_IS_MAPPED
Snapshot that is mapped to a host cannot be deleted
- VOLUME_IS_BOUND
Volume is bound to a ALU
Troubleshooting: Unbound the volume from the ALU
- SNAPSHOT_HAS_ACTIVE_SYNC_JOB
Snapshot is currently a target of an active sync job
Troubleshooting: Please wait for sync job to complete

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
Snapshot Group name does not exist.

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
Snapshot Group name does not exist.

- **MAX_VOLUMES_REACHED**
Maximum number of volumes already defined
- **MAX_COMPRESSED_VOLUMES_REACHED**
Maximum number of compressed volumes already defined
- **DOMAIN_MAX_VOLUMES_REACHED**
The domain exceeds the maximum allowed number of volumes.
- **SNAPSHOT_GROUP_NAME_EXISTS**
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.

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 executed 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**
Snapshot is currently a target of an active sync job
Troubleshooting: Please wait for sync job to complete
- **SNAPSHOT_GROUP_IS_FORMATTED**
Snapshot group is formatted.
- **ELCS_GROUP_CANNOT_BE_FORMATTED**
The snapshot group is an ELCS and cannot be formatted.
- **MAX_VOLUMES_REACHED**
Maximum number of volumes already defined
- **SNAPSHOT_GROUP_IS_INTERNAL**
Internal snapshots cannot be mapped, modified in any way or deleted.
- **VOLUME_IS_NOT_A_SNAPSHOT**

Operation is permitted only on snapshots

- **SNAPSHOT_GROUP_BAD_NAME**
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.

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 in 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

Return codes

- **CONS_GROUP_DOES_NOT_EXIST**
Consistency Group does not exist.
- **SNAPSHOT_GROUP_BAD_NAME**
Snapshot Group name does not exist.

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
Snapshot Group name does not exist.
- SNAPSHOT_GROUP_IS_INTERNAL
Internal snapshots cannot be mapped, modified in any way or deleted.

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

The command is inapplicable for a snapshot group of a mirrored snapshot group.

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**
Snapshot Group name does not exist.
- **SNAPSHOT_GROUP_NAME_EXISTS**
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.

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**
Snapshot Group name does not exist.
- **CONS_GROUP_MISMATCH**
Snapshot Group does not match Consistency Group volumes.
- **NOT_ENOUGH_SPACE**
No space to allocate volume

- **VOLUME_HAS_MIRROR**
Mirror is defined for this volume
- **CONS_GROUP_HAS_MIRROR**
Consistency Group has mirroring defined for it.
- **VOLUME_LOCKED**
Volume is locked
- **TARGET_SNAPSHOT_GROUP_BAD_NAME**
Target Snapshot Group name does not exist.
- **SNAPSHOT_GROUP_MISMATCH**
Snapshot Group does not match target Snapshot Group.
- **TARGET_SNAPSHOT_GROUP_SAME_AS_SOURCE**
Target Snapshot Group is the same as Snapshot Group.
- **VOLUME_HAS_TRANSFORM**
Operation is not supported for volumes which are compressed/decompressed
- **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.
- **COMPRESSION_VOLUME_IS_OFFLINE**
Operation cannot be completed because the requested compressed volume is offline

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**
Snapshot Group name does not exist.
- **SNAPSHOT_GROUP_IS_INTERNAL**
Internal snapshots cannot be mapped, modified in any way 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.

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
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
Snapshot Group name does not exist.

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

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.
- [Listing thin pool compression savings](#)
Use the **pool_compression_saving_list** command to list compression savings.
- [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
gp	Object name	Name of the Grouped Pool to which to move.	N	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.

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**
Consistency Group name does not exist.
 - **POOL_DOES_NOT_EXIST**
Storage Pool does not exist
 - **NOT_ENOUGH_SPACE**
No space to allocate volume
 - **NOT_ENOUGH_HARD_SPACE**
No space to allocate for volume's current usage
 - **NOT_ENOUGH_SNAPSHOT_SPACE**
Snapshot usage will exceed snapshot limit
 - **DOMAIN_MAX_VOLUMES_REACHED**
The domain exceeds the maximum allowed number of volumes.
 - **MAX_VOLUMES_REACHED**
Maximum number of volumes already defined
 - **DOMAIN_MAX_CONS_GROUPS_REACHED**
The domain exceeds the maximum allowed number of consistency groups.
 - **MAX_CONS_GROUPS_REACHED**
Maximum number of Consistency Groups already defined.
 - **DOMAIN_MAX_MIRRORS_REACHED**
The domain exceeds the maximum allowed number of mirrors.
 - **MAX_MIRRORS_REACHED**
Maximum number of mirrors already defined
 - **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 that is 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 entered.
 - **GROUPED_POOL_DOES_NOT_EXIST**
Grouped Pool does not exist.
 - **CONS_GROUP_REQUIRES_DESTINATION_GROUPED_POOL**
A destination Grouped Pool must be entered.
 - **CANNOT_MOVE_CONS_GROUP_TO_A_GP_WITH_NO_META_POOL**
Cannot move Consistency Group to a grouped pool with no meta pool.
 - **BOUND_ALUS_NOT_IN_DESTINATION**
An ALU that is bound to a volume in the pool is not associated with the target domain.
 - **MAX_DMS_REACHED**
Maximum number of remote volumes (mirror/migration) is already defined
- Troubleshooting:** Delete unnecessary Data Migration objects
- **DOMAIN_MAX_DMS_REACHED**
The domain exceeds the maximum allowed number of data migrations.
 - **COMPRESSION_REQUIRES_THIN_PROVISIONED_POOL**
Compressed volumes require thin provisioned pools

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> > ] [ compress=<yes|no> ]
```

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"
compress	Boolean	Specify if new volumes created in this pools are compressed by default or not.	N	no

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 executed 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
Storage Pool does not exist
- PERF_CLASS_BAD_NAME
Performance Class does not exist
- POOL_ALREADY_IN_PERF_CLASS
Pool *pool name* 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 event code

Troubleshooting: Consult the manual for the list of event codes

- EVENT_DOES_NOT_HAVE_THRESHOLDS
Event does not have thresholds

- **EVENT_THRESHOLD_IS_ILLEGAL**
Illegal value for event threshold
- **Troubleshooting:** Event threshold values must be monotonic
- **PERF_CLASS_ASSOCIATED_WITH_HOSTS**
Performance Class *Performance Class* is already in use by host.
- **COMPRESSION_REQUIRES_THIN_PROVISIONED_POOL**
Compressed volumes require thin provisioned pools
- **COMPRESSION_IS_NOT_ACTIVE**
Compression was not enabled.

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 protected level of 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 protected level of 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 protected level of 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**
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 | < hard_size=GB soft_size=GB > >  
snapshot_size=GB [ lock_behavior=<read_only|no_io> ]  
[ perf_class=perfClassName ] [ domain=DomainName ] [ compress=<yes|no> ]
```

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
hard_size	Positive integer	Hard size of the Storage Pool (actual physical capacity).	N	N/A
soft_size	Positive integer	Soft size of the Storage Pool. Soft_size cannot be less than hard_size.	N	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
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
compress	Boolean	Specify if by default new volumes created in this pools are compressed or not	N	no

This command creates a storage pool. The name of the storage pool must be unique in the system. Upon creation, the storage pool is initially empty and does not contain volumes.

The 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 16×2^{30} bytes. The size parameter is used when hard size and soft size are identical (no thin provisioning). If only the size is specified, then hard_size and soft_size are identical. Otherwise, a storage pool with thin provisioning is created.

The created pool has the following values:

- `create_last_consistent_snapshot=yes` - meaning the volumes of this pool can be mirrored.
- `protected_snapshot_priority=2` - managing the way last_consistent snapshots are preserved.

When a storage pool is defined, the new storage pool's capacity is reduced from the system's free space (hard and soft). This operation fails if the system hard or soft free space does not have free capacity of at least the size of the new storage pool. The sum of the capacities of all the storage pools in the system, together with the free space, is always equal to the entire system capacity available for the user.

The system allows for the assignment of the entire available capacity to user created storage pools, while leaving the free space at zero size.

Both hard and soft sizes are subtracted from the free hard/soft space.

For thin provisioned storage pools, the `lock_behavior` parameter sets whether and how the pool is locked upon space depletion. The pool can be locked for write, or for both read and write.

Example:

```
pool_create pool=DBPool size=1000
```

Output:

```
Command completed successfully.
```

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

- **POOL_SNAPSHOT_SIZE_TOO_SMALL**
Snapshot size is very small. It will only allow snapshots of volumes which do no change. All other snapshots will be deleted immediately. Are you sure?

Return codes

- **POOL_NAME_EXISTS**
Storage Pool name already assigned to another Storage Pool
- **PERF_CLASS_BAD_NAME**
Performance Class does not exist
- **MAX_POOLS_REACHED**
Maximum number of Storage Pools already defined
- **NO_HARD_SPACE**
The system does not have enough free hard space for the requested Storage Pool hard size
- **NO_SOFT_SPACE**
The system does not have enough free soft space for the requested Storage Pool soft size
- **NO_SPACE**
The system does not have enough free space for the requested Storage Pool size
- **SOFT_SIZE_SMALLER_THAN_HARD_SIZE**
Soft size must be equal or larger than hard size
- **HARD_SIZE_SMALLER_THAN_SNAPSHOT_SIZE**
Snapshot size must be equal or smaller than hard size
- **REACHED_POOL_MAX_HARD_CAPACITY**
Reached max pool hard capacity
- **DOMAIN_DOESNT_EXIST**
Domain does not exist.
- **USER_ASSOCIATED_TO_MORE_THAN_ONE_DOMAIN**
As the user that runs this command is attached to more than one domain, it is not clear in which domain the pool is created. Run the command again and specify a domain.
- **NO_FREE_HARD_CAPACITY_IN_DOMAIN**
There is not enough free hard space in the domain.
- **NO_FREE_SOFT_CAPACITY_IN_DOMAIN**
There is not enough free soft space in the domain.
- **NO_FREE_CAPACITY_IN_DOMAIN**
There is not enough free space in the domain.
- **DOMAIN_MAX_POOLS_REACHED**
The maximum number of domain pools was reached.
- **PERF_CLASS_ASSOCIATED_WITH_HOSTS**
Performance Class *Performance Class* is already in use by host.
- **COMPRESSION_REQUIRES_THIN_PROVISIONED_POOL**
Compressed volumes require thin provisioned pools
- **COMPRESSION_IS_NOT_ACTIVE**
Compression was not enabled.

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
Storage Pool does not exist
- POOL_HAS_CG
Storage Pool has Consistency Groups defined
- POOL_IN_USE
Storage Pool has volumes allocated in it
- 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.
- POOL_BELONGS_TO_A_GROUPED_POOL
Pool belongs to a Grouped Pool.

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.
gp	Object name	Name of a Group Pool.	N	All pools.
managed	Boolean	Determines whether to show unmanaged pools (no), managed (yes), or both (all).	N	No
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
soft_size	Size (GB)	2
soft_size_MiB	Size (MiB)	N/A

Field ID	Field output	Default position
hard_size	Hard Size (GB)	6
hard_size_MiB	Hard Size (MiB)	N/A
snapshot_size	Snap Size (GB)	4
snapshot_size_MiB	Snap Size (MiB)	N/A
total_volume_size	Soft Vols (GB)	3
total_volume_size_MiB	Soft Vols (MiB)	N/A
empty_space_soft	Soft Empty (GB)	5
empty_space_soft_MiB	Soft Empty (MiB)	N/A
empty_space_hard	Hard Empty (GB)	10
empty_space_hard_MiB	Hard Empty (MiB)	N/A
used_by_volumes	Hard Vols (GB)	7
used_by_volumes_MiB	Hard Vols (MiB)	N/A
used_by_snapshots	Hard Snaps (GB)	9
used_by_snapshots_MiB	Hard 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
compress	Create Compressed Volumes	13
managed	Managed	N/A
perf_class	Perf Class Name	11
domain	Domain	12

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Listing thin pool compression savings

Use the **pool_compression_saving_list** command to list compression savings.

```
pool_compression_saving_list [ pool=PoolName ] [ domain=DomainName ]
```

Parameters

Name	Type	Description	Mandatory	Default
pool	Object name	Name of a specific thin storage pool to be listed.	N	All thin pools.
domain	Object name	Name of a specific domain to be listed.	N	All Domains

If no parameter is specified, this command lists the savings for all of the accessible pools.

If a domain is specified, then the totals for all volumes in the specified domain are displayed. If the pool parameter is specified, the command lists the totals for the specific pool: the sum of all of the compressed volumes in the given pool.

Example:

```
xcli -u -c XIV1 pool_compression_saving_list pool
```

Output:

Command completed successfully

Field ID	Field output	Default position
name	Name	1
size	Size (GB)	2
size_MiB	Size (MiB)	N/A
num_compressed_volumes	Compressed Volumes	3
saving	Compression Estimated Saving (GB)	4
saving_MiB	Compressed Saving (MiB)	N/A
creator	Creator	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

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	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- POOL_DOES_NOT_EXIST
Storage Pool does not exist
- POOL_NAME_EXISTS
Storage Pool name already 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
hard_size	Positive integer	Hard size of the storage pool (actual physical capacity).	N	N/A

Name	Type	Description	Mandatory	Default
soft_size	Positive integer	Soft size of the storage pool (maximal size of capacity seen by the hosts, used for thin provisioning).	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 16×2^{30} bytes.

Capacity accounting is performed in respect to the free space.

You can either specify both hard_size and soft_size or to only specify size (which specifies 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 in size.
- When decreasing a storage pool size, the command succeeds only if the storage pool itself holds enough free capacity to allow such a reduction in size.
- If the new size equals the current size, the command succeeds without changes to the storage pool.

This command fails if either the current storage pool's size (hard or soft) cannot be decreased or if free space (hard or soft) cannot be decreased.

Example:

```
pool_resize pool=DBPool size=1300
```

Output:

```
Command executed 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_SNAPSHOT_SIZE_TOO_SMALL**
Snapshot size is very small. It will only allow snapshots of volumes which do no change. All other snapshots will be deleted immediately. Are you sure?

Return codes

- **POOL_DOES_NOT_EXIST**
Storage Pool does not exist
- **NO_SOFT_SPACE**
The system does not have enough free soft space for the requested Storage Pool soft size
- **SOFT_SIZE_SMALLER_THAN_HARD_SIZE**
Soft size must be equal or larger than hard size
- **HARD_SIZE_SMALLER_THAN_SNAPSHOT_SIZE**
Snapshot size must be equal or smaller than hard size
- **POOL_SOFT_TOO_SMALL**
Requested soft size is smaller than the sum of sizes of volumes in the Storage Pool
- **POOL_TOO_SMALL**
Storage Pool usage exceeds requested size
- **POOL_HARD_TOO_SMALL**
Storage Pool usage exceeds requested hard size
- **NO_SPACE**
The system does not have enough free space for the requested Storage Pool size
- **NO_HARD_SPACE**
The system does not have enough free hard space for the requested Storage Pool hard size
- **REACHED_POOL_MAX_HARD_CAPACITY**
Reached max pool hard capacity
- **POOL_MUST_BE_THIN**
This pool has a golden snapshot and therefore must be thin provisioned.

- 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.
- GROUPED_POOL_THIN_MUST_BE_THIN
Thin pool must be thin provisioned.
- NO_FREE_HARD_CAPACITY_IN_DOMAIN
There is not enough free hard space in the domain.
- NO_FREE_SOFT_CAPACITY_IN_DOMAIN
There is not enough free soft space in the domain.
- NO_FREE_CAPACITY_IN_DOMAIN
There is not enough free space in the domain.
- COMPRESSION_REQUIRES_THIN_PROVISIONED_POOL
Compressed volumes require thin provisioned pools

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 <i>true</i> , 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
Volume name does not exist
- POOL_DOES_NOT_EXIST
Storage Pool does not exist
- NOT_ENOUGH_SPACE
No space to allocate volume
- NOT_ENOUGH_HARD_SPACE
No space to allocate for volume's current usage

- **VOLUME_IS_SNAPSHOT**
Operation is not permitted on snapshots
 - **VOLUME_HAS_OLVM**
IBM Hyper-Scale Mobility relation is defined for this volume
 - **VOLUME_BELONGS_TO_CG**
Volume belongs to a Consistency Group
 - **NOT_ENOUGH_SNAPSHOT_SPACE**
Snapshot usage will exceed 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.
 - **CANNOT_MOVE_TO_THICK_POOL_VOLUME_HAS_GOLDEN_SNAPSHOTS**
Volume has golden snapshots and therefore cannot be moved to a thick pool.
 - **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 that is 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**
Maximum number of mirrors already defined
 - **MAX_DMS_REACHED**
Maximum number of remote volumes (mirror/migration) is already defined
- Troubleshooting:** Delete unnecessary Data Migration objects
- **MAX_VOLUMES_REACHED**
Maximum number of volumes already defined
 - **BOUND_ALUS_NOT_IN_DESTINATION**
An ALU that is bound to a volume in the pool is not associated with the target domain.
 - **COMPRESSION_REQUIRES_THIN_PROVISIONED_POOL**
Compressed volumes require thin provisioned pools
 - **VOLUME_HAS_TRANSFORM**
Operation is not supported for volumes which are compressed/decompressed

System management commands

This section describes the command-line interface (CLI) for system management.

- [Displaying current consumed capacity of the system](#)
Use the `cod_list` to display current consumed capacity of the system.
- [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.
- [Displaying 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 the system power management data](#)
Use the **system_power_management_get** command to display system runtime, power consumption, system alert, and number of PSUs without power input.
- [Displaying information about effective and physical capacity](#)
Use the **system_capacity_list** command to display information about the system's current hard, soft, and allocated soft capacities.
- [Setting the system's soft capacity](#)
Use the **system_soft_capacity_set** to set the system's soft capacity.
- [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 values of maintenance module parameters](#)
Use the **mm_config_get** command to display the values of maintenance module 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.
- [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 current consumed capacity of the system

Use the **cod_list** to display current consumed capacity of the system.

```
cod_list [ name=Name ]
```

Parameters

Name	Type	Description	Mandatory	Default
name	String	Name of parameter to display.	N	All parameters
Field ID	Field name	Default position		
name	Name	1		
value	Value	2		

This command displays current consumed capacity of a given system.

Example:

```
cod_list
```

Output:

```

Name      Value
-----
consumed_capacity  1039
date            2009-05-27
dst             yes

```

```

machine_model      A14
machine_serial_number MN00013
machine_type       2810
system_id          13
system_name        XIV MN00013a
time               10:13:31
timezone           Asia/Jerusalem

```

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

- CONF_SERVER_UNREACHABLE
Configuration server unreachable
- UNRECOGNIZED_CONFIG_PARAMETER
Unrecognized configuration parameter: '*name*'.

Troubleshooting: Use a valid configuration parameter as an input.

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.
- support_center_port_type -
- system_id - Unique system identifier (equivalent to a serial number).
- system_name

Example:

```
config_get
```

Output:

Name	Value
-----	-----

```
email_sender_address    support@ibm.com
email_reply_to_address  storage@ibm.com
dns_primary             10.0.0.10
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
Technicians	Allowed

Return codes

- CONF_SERVER_UNREACHABLE
Configuration server unreachable
- UNRECOGNIZED_CONFIG_PARAMETER
Unrecognized configuration parameter: 'name'.

Troubleshooting: Use a valid configuration parameter as an input.

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.

Example:

```
config_set name=dns_secondary value=10.0.0.119
```

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	Allowed

Return codes

- UNRECOGNIZED_CONFIG_PARAMETER
Unrecognized configuration parameter: '*name*'.
- READ_ONLY_CONFIG_PARAMETER
Configuration parameter: '*name*' is read-only.
- IPV4_NOT_CONFIGURED
IPv4 address is not configured on management interface
- RULE_WITH_SNMP_DEST_EXISTS
Can not set snmp_type to none. There is a rule that contains SNMP destination.

Troubleshooting: Use a valid configuration parameter as an input.

Troubleshooting: You cannot modify read-only parameters.

Troubleshooting: Define IPv4 address for management before disabling IPv6

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	212.143.102.243	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 with short output (default for command) - Displays the command name and short description.
- Command with full output (default when used in XIV-internal mode) - Displays the command name, short description, syntax, list of parameters and their description, types and default values. If output is table, displays all possible table columns.

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

Displaying 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
Read-only users	Allowed
Technicians	Allowed

Shutting down the system

Use the **shutdown** command to shut down the system.

```
shutdown [ emergency=<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

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.

NOTE: USING THIS OPTION MAY CAUSE DATA LOSS.

Example:

```
shutdown
```

Output:

```
Command executed 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 support

- CANNOT_WRITE_TO_KEY_REPOSITORY
Failed writing keys to the key repository.

Troubleshooting: Contact support.

- NO_LIVE_KEYSERVER_GATEWAY_NODE
There is no live key server gateway node on the system.

Troubleshooting: Please restart the key server gateway node and try again.

- NO_MASTER_KEYSERVER_DEFINED
There is no master key server defined on the system.

Troubleshooting: Please 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.

Listing the operational state

Use the **state_list** command to display the current operational state of the system.

state_list

Field ID	Field output	Default position
category	Category	1
value	Value	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

Displaying the system power management data

Use the **system_power_management_get** command to display system runtime, power consumption, system alert, and number of PSUs without power input.

system_power_management_get

This command lists the following:

- *system_runtime* - The system remaining runtime.
- *stem_required_runtime* - The minimum required runtime.
- *system_power_consumption* - The system power consumption.
- *num_of_psu_no_ac_in* - The number of PSUs in the system with no input power.
- *runtime_reference* - The runtime reference Actual/Predictive.

Example:

system_power_management_get

Output:

Runtime	Min Required Runtime	Power Consumption	PSUs with no input power	Runtime Reference
11	5	5979	0	ACTUAL

Field ID	Field output	Default position
system_runtime	Runtime	1
system_required_runtime	Min Required Runtime	2
system_power_consumption	Power Consumption	3
num_of_psu_no_ac_in	PSUs with no input power	4
runtime_reference	Runtime Reference	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

Displaying information about effective and physical capacity

Use the **system_capacity_list** command to display information about the system's current hard, soft, and allocated soft capacities.

system_capacity_list

The output includes the following fields:

Field ID	Field output	Default position
soft	Soft	1
hard	Hard	2
max_pool_size	Max_Pool_Size	3
free_hard	Free Hard	4
free_soft	Free Soft	5
spare_modules	Spare Modules	6
spare_disks	Spare Disks	7
target_spare_modules	Target Spare Modules	8
target_spare_disks	Target Spare Disks	9
soft_MiB	Soft (MiB)	N/A
hard_MiB	Hard (MiB)	N/A
free_hard_MiB	Free Hard (MiB)	N/A
free_soft_MiB	Free Soft (MiB)	N/A
capacity_limit_percentage	Capacity Limit (%)	10

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Setting the system's soft capacity

Use the **system_soft_capacity_set** to set the system's soft capacity.

system_soft_capacity_set soft_size=SizeGB

Parameters

Name	Type	Description	Mandatory
soft_size	Positive integer	The system's new soft size. This is the user space available to applications with thin provisioning. Maximum allowed soft size is 2PB. Minimum allowed soft size is system_hard_capacity.	Y

The soft capacity size of the system can be set to up to 2PB, and as low as the maximum size between the currently allocated soft capacity and the system's hard capacity (whichever is greater). The current hard, soft, and/or allocated soft capacity can be retrieved using the **system_capacity_list** command.

Example:

system_soft_capacity_set soft_size=250000

Output:

None

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_SYSTEM_SOFT_CAPACITY
Are you sure you want to update system soft capacity?

Return codes

- SOFT_CAPACITY_TOO_HIGH
The given soft capacity is too high
Troubleshooting: Check documentation for legal soft capacity values
- SOFT_SIZE_SMALLER_THAN_HARD_SIZE
Soft size must be equal or larger than hard size
- SOFT_SIZE_DECREASE_FAILED
System soft size cannot be less than size already allocated

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 executed 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**
System time was not changed
- **Troubleshooting:** Please try again.
- **BAD_TIMESTAMP**
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	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
Configuration server unreachable
- UNRECOGNIZED_CONFIG_PARAMETER
Unrecognized configuration parameter: '*name*'.

Troubleshooting: Use a valid configuration parameter as an input.

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.PID`

Example:

```
vpd_config_set name= value=
```

Output:

Command executed 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

- `READ_ONLY_CONFIG_PARAMETER`
Configuration parameter: '*name*' is read-only.
Troubleshooting: You cannot modify read-only parameters.
- `UNRECOGNIZED_CONFIG_PARAMETER`
Unrecognized configuration parameter: '*name*'.
Troubleshooting: Use a valid configuration parameter as an input.

Displaying values of maintenance module parameters

Use the `mm_config_get` command to display the values of maintenance module parameters.

```
mm_config_get [ name=Name ]
```

Parameters

Name	Type	Description	Mandatory	Default

Name	Type	Description	Mandatory	Default
name	String	Name of the parameter to print.	N	All of the parameters.
ID	Name	Default Position		
name	Name	1		
value	Value	2		

Example:

```
mm_config_get name=mm_mutable_info.should_run_package_daemon
```

Output:

Name	Value
-----	-----
mm_mutable_info.should_run_package_daemon	yes

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

- CONF_SERVER_UNREACHABLE
Configuration server unreachable
- UNRECOGNIZED_CONFIG_PARAMETER
Unrecognized configuration parameter: '*name*'.

Troubleshooting: Use a valid configuration parameter as an input.

Displaying the system's MIB file

Use the **mib_get** command to display the system's MIB file.

mib_get

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 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
Application administrator	Allowed

User Category	Permission
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Return codes

- CANNOT_READ_FROM_FILE
Cannot read from file '*Filename*'

Troubleshooting: Contact 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.	Y	N/A
approver_name	String	The approver's name.	N	none

Example:

```
elicense_accept version approver_name
```

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.

Example:

```
xccli -u -c XIV1 audit_enable
```

Output:

Command executed 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

- **AUDIT_ALREADY_ENABLED**
Command auditing 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 executed 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

- **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
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	Disallowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- **AUDIT_PRIMARY_SAME_AS_SECONDARY**
The same audit server is defined as both primary and secondary.

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
198.51.100.42	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

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 executed 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
CIM port 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 executed 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
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.

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
Target name does not exist
- TARGET_INVALID_CONNECTION_DURATION_THRESHOLD
Target connection duration threshold should be in [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

User Category	Permission
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- TARGET_BAD_NAME
Target name does not exist
- TARGET_ILLEGAL_RATE_VALUES
max init rate should be smaller or equal to max sync job rate. max sync job rate should not be greater than max resync 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
Target name does not exist
- CONNECTIVITY_NOT_DEFINED
Remote port is not connected through this local port
- COMPONENT_IS_NOT_AN_FC_PORT
Component must specify an FC port
- COMMAND_NOT_ALLOWED_ON_MANAGEMENT_OR_VPN_INTERFACE
Operation is not allowed on Management or VPN IP Interface
- IPINTERFACE_DOES_NOT_EXIST
IP Interface name does not exist
- TARGET_PORT_BAD_ADDRESS
Remote port address is illegal or does not belong to the remote target
- BAD_LOCAL_IP_PORT
An 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> ]
```

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	Reserved	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
Target name does not exist
- COMMAND_NOT_ALLOWED_ON_MANAGEMENT_OR_VPN_INTERFACE
Operation is not allowed on Management or VPN IP Interface
- CONNECTIVITY_NOT_DEFINED
Remote port is not connected through this local port
- COMPONENT_IS_NOT_AN_FC_PORT
Component must specify an FC port
- TARGET_PORT_BAD_ADDRESS
Remote port address is illegal or does not belong to the remote target
- BAD_LOCAL_IP_PORT
An ID of a local IP port must be specified
- IPINTERFACE_DOES_NOT_EXIST
IP Interface name does not exist
- TARGET_HAS_OLVM_RELATIONSHIP
Target has an IBM Hyper-Scale Mobility relationship - cannot be deactivate 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 > |  
< 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	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 de-activate (target_connectivity_deactivate) and then re-activate (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.

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
Target name does not exist
- CONN_EXISTS
Remote port is already connected through this local port
- MAX_CONNECTIONS_REACHED
Maximum number of connections already defined
- MAX_ISCSI_CONNECTIONS_PER_MODULE_REACHED
Maximal number of iSCSI connectivities already defined for that module.
- COMPONENT_IS_NOT_AN_FC_PORT
Component must specify an FC port
- COMPONENT_IS_NOT_FC_INITIATOR_PORT
Component must specify FC initiator port
- BAD_LOCAL_IP_PORT
An ID of a local IP port must be specified
- COMMAND_NOT_ALLOWED_ON_MANAGEMENT_OR_VPN_INTERFACE
Operation is not allowed on Management or VPN IP Interface
- IPINTERFACE_DOES_NOT_EXIST
IP Interface name does not exist
- TARGET_PORT_BAD_ADDRESS
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> ]
```

Parameters

Name	Type	Description	Mandatory	Default
target	Object name	Remote target of the connectivity definition.	Y	N/A

Name	Type	Description	Mandatory	Default
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	Reserved	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
Target name does not exist
- COMMAND_NOT_ALLOWED_ON_MANAGEMENT_OR_VPN_INTERFACE
Operation is not allowed on Management or VPN IP Interface
- CONNECTIVITY_NOT_DEFINED
Remote port is not connected through this local port
- COMPONENT_IS_NOT_AN_FC_PORT
Component must specify an FC port
- TARGET_PORT_BAD_ADDRESS
Remote port address is illegal or does not belong to the remote target
- BAD_LOCAL_IP_PORT
An ID of a local IP port must be specified
- IPINTERFACE_DOES_NOT_EXIST
IP Interface name does not exist
- TARGET_HAS_OLVM_RELATIONSHIP
Target has an IBM Hyper-Scale Mobility relationship - cannot be deactivate 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 ]  
[ uses_4k_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 the system_id parameter on the remote system (see Displaying the values of configuration parameters).	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
uses_4k_sector s	Boolean	Optimize the Asynchronous mirror data transfer for remote targets with 4KB sector size, such as 6TB or larger disks.	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.

Larger capacity disks, such as 6TB and 10TB, use 4KB sector size and therefore may suffer performance degradation as a direct result of an existing performance optimizations added for the benefit of smaller sector disks. The uses_4k_sectors parameter allows the user to prevent the performance optimization in the cases where it may prove harmful. The disks affected by the performance optimization are large capacity disks (6TB and higher) inside the remote target system.

Example:

```
target_define target=Nextra2 protocol=FC
```

Output:

```
Command executed 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 slave system. Are you sure the remote system is not already defined as a target?

Return codes

- **MAX_TARGETS_REACHED**
Maximum number of targets already defined
- **TARGET_NAME_EXISTS**
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**
FC Target does not have an iscsi_name
- **TARGET_BAD_SCSI_TYPE**
Target SCSI type does not exist
- **DOMAIN_DOESNT_EXIST**
Domain does not exist.

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	Reserved	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 executed 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**
Target name does not exist
- **TARGET_HAS_PORTS**
Ports are defined for this target
- **TARGET_HAS_ASSOCIATIONS**
Remote volumes are defined on this target
- **TARGET_HAS_OLVM_RELATIONSHIP**
Target has an IBM Hyper-Scale Mobility relationship - cannot be deactivate 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	Connected	3
xiv_target	XIV Target	N/A
iscsi_name	iSCSI Name	N/A
system_id	System ID	N/A
num_ports	Number of Ports	N/A
creator	Creator	N/A
max_initialization_rate	Max Initialization Rate	4
max_resync_rate	Max Resync Rate	5
max_syncjob_rate	Max Syncjob Rate	6
connectivity_lost_event_threshold	Connection Threshold	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 executed successfully
```

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

- TARGET_BAD_NAME
Target name does not exist
- TARGET_BAD_TYPE
Target machine is not 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=Nexttra2 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
Remote port address is illegal or does not belong to the remote target
- TARGET_BAD_PORT_STATE
Port is already in requested activation state
- TARGET_BAD_NAME
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

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
Target name does not exist
- PORT_EXISTS
Port is already defined
- MAX_PORTS_REACHED
Maximum number of ports already defined in the system
- TARGET_PORT_BAD_ADDRESS
Remote port address is illegal or does not belong to the remote target
- ISCSI_HOST_ILLEGAL_PORT_NAME
Port name for iSCSI Host is illegal

Troubleshooting: Port names for iSCSI Hosts must contain only printable characters.

- HOST_PORT_EXISTS
Host with this port ID 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> ]
```

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	Reserved	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

User Category	Permission
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- TARGET_BAD_NAME
Target name does not exist
- TARGET_PORT_BAD_ADDRESS
Remote port address is illegal or does not belong to the remote target
- TARGET_BAD_PORT_STATE
Port is already in requested activation state
- TARGET_HAS_OLVM_RELATIONSHIP
Target has an IBM Hyper-Scale Mobility relationship - cannot be deactivate 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 >
```

Parameters

Name	Type	Description	Mandatory
target	Object name	Remote target from which the port is that is deleted.	Y
ipaddress	N/A	IP address of the port (for iSCSI targets only).	N
fcaddress	N/A	FC address of the remote port (for FC targets only).	N

Example:

```
target_port_delete
target=Nexttra2
fcaddress=10:00:00:17:38:27:ec:11
```

Output:

```
Command executed 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
Target name does not exist
- TARGET_PORT_BAD_ADDRESS
Remote port address is illegal or does not belong to the remote target
- TARGET_PORT_HAS_CONNECTIVITY
Port has connectivity defined to it

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 executed 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
Target name does not exist
- TARGET_NAME_EXISTS
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_4k_sectors=<yes|no> ]
```


Parameters

Name	Type	Description	Mandatory	
target	Object name	Target to be updated.	Y	
system_id	String	ID of the remote system. Should be the same as the output of Displaying the values of configuration parameters of the <code>system_id</code> variable on the remote system.	Y	
uses_4k_sectors	Boolean	Optimize the Asynchronous mirror data transfer for remote targets with 4KB sector size, such as 6TB or larger disks.	N	No

This command changes the system ID of the remote target.

Larger capacity disks, such as 6TB and 10TB, use 4KB sector size and therefore may suffer performance degradation as a direct result of an existing performance optimizations added for the benefit of smaller sector disks. The `uses_4k_sectors` parameter allows the user to prevent the performance optimization in the cases where it may prove harmful. The disks affected by the performance optimization are large capacity disks (6TB and higher) inside the remote target system.

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`
Target name does not exist

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 the 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 roles of a mirrored volume](#)
Use the `mirror_change_role` command to change the role of a local mirroring peer between master and slave.
- [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 canceled.	N	N/A
cg	Object name	Name of the (local) master consistency group whose non-started snapshot mirrors should be canceled.	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 canceled. 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 canceled.

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
Storage integration administrator	Allowed	N/A
Application administrator	Conditionally Allowed	The volume is mapped to a host or a cluster associated with the user.
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
Volume name does not exist
- VOLUME_NO_MIRROR
Local volume does not have remote mirroring definitions
- CONS_GROUP_BAD_NAME
Consistency Group name does not exist.
- CONS_GROUP_NO_MIRROR
Local Consistency Group does not have remote mirroring definitions

- LOCAL_PEER_IS_NOT_MASTER
Local peer is not the master
- VOLUME_BELONGS_TO_MIRRORED_CONS_GROUP
Volume mirror is part of Consistency Group mirror.
- MIRROR_RETRY_OPERATION
There is an operation in progress on this mirror , please try again your request in a few seconds

Troubleshooting: Please try again the command in a few seconds

- TARGET_BAD_NAME
Target name does not exist
- VOLUME_HAS_MULTIPLE_MIRRORS
volume has multiple mirrors, operation not allowed or target must be specified
- VOLUME_TARGET_MISMATCH
Volume and target do not match
- CONS_GROUP_BAD_TARGET
Target name does not match the Consistency Group
- MIRROR_IS_STANDBY
Mirror is marked as standby

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 ]
slave_name=SnapshotName [ slave_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 mirror name. Mandatory if 2 mirrors are defined on the volume.	N	[none]
name	Object name	The name of the new snapshot.	Y	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.	Y	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

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.
Security administrator	Disallowed	N/A
Read-only users	Disallowed	N/A
Technicians	Disallowed	N/A

Return codes

- **CONS_GROUP_MISMATCH**
Snapshot Group does not match Consistency Group volumes.
- **CONS_GROUP_EMPTY**
Operation is not allowed on an empty Consistency Group.
- **CONS_GROUP_BAD_NAME**
Consistency Group name does not exist.
- **CONS_GROUP_NO_MIRROR**
Local Consistency Group does not have remote mirroring definitions
- **LOCAL_PEER_IS_NOT_MASTER**
Local peer is not the master
- **MIRROR_IS_NOT_SYNCHRONIZED**
Mirror is not synchronized
- **MIRROR_RETRY_OPERATION**
There is an operation in progress on this mirror , please try again your request in a few seconds

Troubleshooting: Please try again the command in a few seconds

- **MIRROR_IS_NON_OPERATIONAL**
Mirror is non-operational
- **MAX_VOLUMES_REACHED**
Maximum number of volumes already defined
- **DOMAIN_MAX_VOLUMES_REACHED**
The domain exceeds the maximum allowed number of volumes.
- **OPERATION_NOT_ALLOWED_ON_LOOPBACK**
Requested operation is not allowed on loopback target
- **OVERWRITE_SNAPSHOT_BAD_NAME**
Snapshot name does not exist
- **OVERWRITE_SNAPSHOT_GROUP_DOES_NOT_BELONG_TO_GIVEN_GROUP**
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 to create a snapshot on remote.
- **REMOTE_MAX_VOLUMES_REACHED**
Maximum number of volumes already defined on remote machine
- **REMOTE_MAX_SNAPSHOTS_FOR_VOLUME_REACHED**
Maximal number of snapshots per volume is already reached on a remote whose version is not 10.2.4.
- **REMOTE_VOLUME_IS_MASTER**
Volume on remote machine is currently defined as Master
- **REMOTE_SNAPSHOT_NAME_EXISTS**
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
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
Remote Snapshot Group name has a reserved prefix.
- REMOTE_SNAPSHOT_BAD_PREFIX
Remote snapshot name has a reserved prefix
- SNAPSHOT_HAS_ACTIVE_SYNC_JOB
Snapshot is currently a target of an active sync job

Troubleshooting: Please wait for 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 in any way or deleted.
- SNAPSHOT_GROUP_NAME_EXISTS
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
Snapshot Group name does not exist.
- SNAPSHOT_GROUP_BAD_PREFIX
Snapshot Group name has a reserved prefix.
- SNAPSHOT_IS_PART_OF_SNAPSHOT_GROUP
Snapshot is part of a Snapshot Group
- SYNCHED_SNAPSHOTS_NOT_SUPPORTED_IN_TARGET
Synchronized Snapshot capability is not supported by the Mirror's Target.
- VOLUME_BAD_PREFIX
Volume name has a reserved prefix
- VOLUME_BELONGS_TO_MIRRORED_CONS_GROUP
Volume mirror is part of Consistency Group mirror.
- VOLUME_DATA_MIGRATION_UNSYNCHRONIZED
Data Migration has not completed to this volume
- VOLUME_EXISTS
Volume name already exists
- VOLUME_BAD_NAME
Volume name does not exist
- VOLUME_NO_MIRROR
Local volume does not have remote mirroring definitions
- VOLUME_IS_NOT_CONSISTENT_SLAVE
Operation not allowed on slave volume that is not consistent.
- VOLUME_IS_SNAPSHOT
Operation is not permitted on snapshots
- VOLUME_HAS_TRANSFORM
Operation is not supported for volumes which are compressed/decompressed
- 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.
- CONS_GROUP_BAD_TARGET
Target name does not match the Consistency Group
- TARGET_BAD_NAME

Target name does not exist

- **VOLUME_TARGET_MISMATCH**
Volume and target do not match
- **REMOTE_MIRROR_IS_STANDBY**
remote mirror is marked as standby
- **VOLUME_HAS_MULTIPLE_MIRRORS**
volume has multiple mirrors, operation not allowed or target must be specified
- **XMIRROR_SNAPSHOT_MIRROR_NOT_SUPPORTED**
Volume is part of an xmirror, snapshot mirrors are not supported for xmirrors.
- **MIRROR_IS_STANDBY**
Mirror is marked as standby

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**
Volume name does not exist
- **VOLUME_NO_MIRROR**
Local volume does not have remote mirroring definitions
- **CONS_GROUP_BAD_NAME**
Consistency Group name does not exist.
- **CONS_GROUP_NO_MIRROR**
Local Consistency Group does not have remote mirroring definitions

- **LOCAL_PEER_IS_NOT_MASTER**
Local peer is not the master
 - **MIRROR_CONFIGURATION_ERROR**
Mirror local configuration does not match remote configuration
 - **REMOTE_MAX_VOLUMES_REACHED**
Maximum number of volumes already defined on remote machine
 - **REMOTE_MAX_COMPRESSED_VOLUMES_REACHED**
Maximum number of compressed volumes already defined in the remote
 - **SYNC_ALREADY_ACTIVE**
Synchronization is already active
 - **VOLUME_BELONGS_TO_MIRRORED_CONS_GROUP**
Volume mirror is part of Consistency Group mirror.
 - **MIRROR_CAN_NOT_BE_ACTIVATED**
Mirroring cannot be activated
 - **MIRROR_CONS_GROUP_MEMBERSHIP_MISMATCH**
Mirrored CG contains different volumes on Master and Slave. This problem occurs whenever the `cg_add_vol` or `cg_remove_vol` commands were previously issued and the Master did not receive an acknowledgment from the Slave until the command timed out, or any other unexpected failure.
 - **MIRROR_SIZE_MISMATCH**
Slave volume and Master Volume sizes are different
 - **MIRROR_RETRY_OPERATION**
There is an operation in progress on this mirror , please try again your request in a few seconds
- Troubleshooting:** Please try again the command in a few seconds
- **VOLUME_HAS_MULTIPLE_MIRRORS**
volume has multiple mirrors, operation not allowed or target must be specified
 - **TARGET_BAD_NAME**
Target name does not exist
 - **VOLUME_TARGET_MISMATCH**
Volume and target do not match
 - **CONS_GROUP_BAD_TARGET**
Target name does not match the Consistency Group
 - **VOLUME_TOO_MANY_ACTIVE_MIRRORS**
This command cannot be used if more than one mirror is active on the volume
 - **REMOTE_MIRROR_IS_STANDBY**
remote mirror is marked as standby
 - **REMOTE_DOMAIN_MAX_VOLUMES_REACHED**
Maximum number of volumes already defined on remote machine domain
 - **MIRROR_IS_STANDBY**
Mirror is marked as standby
 - **COMPRESSION_VOLUME_IS_OFFLINE**
Operation cannot be completed because the requested compressed volume is offline

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 executed 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
Volume name does not exist
- CONS_GROUP_NO_MIRROR
Local Consistency Group does not have remote mirroring definitions
- ASYNC_MIRROR_REMOTE_RPO_TOO_SHORT
Specified Remote RPO is too short.
- ASYNC_MIRROR_RPO_TOO_LONG
Specified RPO is too long.
- TARGET_NOT_CONNECTED
There is currently no connection to the target system
- VOLUME_NO_MIRROR
Local volume does not have remote mirroring definitions
- CONS_GROUP_BAD_NAME
Consistency Group name does not exist.
- ASYNC_MIRROR_RPO_TOO_SHORT
Specified RPO is too short.
- VOLUME_BELONGS_TO_MIRRORED_CONS_GROUP
Volume mirror is part of Consistency Group mirror.
- INTERVAL_SHOULD_BE_SHORTER_THAN_RPO
Schedule interval must be shorter than the RPO.
- ASYNC_MIRROR_REMOTE_RPO_TOO_LONG
Specified Remote RPO is too long.
- LOCAL_IS_SLAVE
Local mirror peer is not the master
- SYNC_MIRROR_HAS_NO_RPO
Synchronous Mirror does not have an RPO.
- TARGET_BAD_NAME
Target name does not exist
- VOLUME_HAS_MULTIPLE_MIRRORS
volume has multiple mirrors, operation not allowed or target must be specified
- VOLUME_TARGET_MISMATCH
Volume and target do not match
- CONS_GROUP_BAD_TARGET
Target name does not match the Consistency Group
- MIRROR_IS_STANDBY
Mirror is marked as standby

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/slave 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 executed 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
Volume name does not exist
- VOLUME_NO_MIRROR
Local volume does not have remote mirroring definitions
- CONS_GROUP_BAD_NAME
Consistency Group name does not exist.
- CONS_GROUP_NO_MIRROR
Local Consistency Group does not have remote mirroring definitions
- LOCAL_PEER_IS_NOT_MASTER
Local peer is not the master
- MIRROR_DESIGNATION_NOT_SUPPORTED_BY_TARGET
Mirror role designation is not supported by the Mirror's Target.
- MIRROR_IS_NON_OPERATIONAL
Mirror is non-operational
- VOLUME_BELONGS_TO_MIRRORED_CONS_GROUP
Volume mirror is part of Consistency Group mirror.
- TARGET_BAD_NAME
Target name does not exist
- VOLUME_HAS_MULTIPLE_MIRRORS
volume has multiple mirrors, operation not allowed or target must be specified
- VOLUME_TARGET_MISMATCH
Volume and target do not match
- CONS_GROUP_BAD_TARGET
Target name does not match the Consistency Group
- MIRROR_IS_STANDBY

Mirror is marked as standby

Changing the 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**
Volume name does not exist
- VOLUME_NO_MIRROR**
Local volume does not have remote mirroring definitions
- CONS_GROUP_BAD_NAME**
Consistency Group name does not exist.
- CONS_GROUP_NO_MIRROR**
Local Consistency Group does not have remote mirroring definitions
- REMOTE_VOLUME_IS_MASTER**
Volume on remote machine is currently defined as Master
- REMOTE_CONS_GROUP_IS_MASTER**
Remote Consistency Group is defined as Master
- SCHEDULE_DOES_NOT_EXIST**
Specified Schedule does not exist

- **SYNC_MIRROR_DOES_NOT_USE_SCHEDULE**
Definition of Synchronous Mirror does not require a Schedule object to be specified.
- **TARGET_NOT_CONNECTED**
There is currently no connection to the target system
- **VOLUME_BELONGS_TO_MIRRORED_CONS_GROUP**
Volume mirror is part of Consistency Group mirror.
- **INTERVAL_SHOULD_BE_SHORTER_THAN_RPO**
Schedule interval must be shorter than the RPO.
- **TARGET_BAD_NAME**
Target name does not exist
- **VOLUME_HAS_MULTIPLE_MIRRORS**
volume has multiple mirrors, operation not allowed or target must be specified
- **VOLUME_TARGET_MISMATCH**
Volume and target do not match
- **CONS_GROUP_BAD_TARGET**
Target name does not match the Consistency Group
- **DOMAIN_HAS_NO_ACCESS_TO_SCHEDULE**
Domain has no access to schedule.
- **MIRROR_IS_STANDBY**
Mirror is marked as standby

Changing the roles of a mirrored volume

Use the **mirror_change_role** command to change the role of a local mirroring peer between master and slave.

```
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	CG 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.

When the command is applied to the master:

- The command can be issued only if the activation state is Standby.
- The command cannot be issued during the initialization phase.

Changing the roles in synchronous mirroring:

- When applied on the master:
 - All changes made to the master since the last time the peers were synchronized will be reverted to their original value. The master ceases serving host requests, and is set to accept replication from the other peer as a slave. If the command is issued during link unavailability, a most_updated snapshot of the peer will be taken to capture the most recent changes that have not yet been replicated to the other peer.
 - A warning is displayed: **Are you sure to change master to slave?"**
 - An event is generated
 - The master ceases accepting host requests
 - Unsynchronized data at the demoted master is recorded in most updated snapshot
 - The demoted master reverts to the last replicated snapshot
 - Completion of process is recorded in the log
- When applied on the 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 slave volume has a last consistent snapshot, it means that the mirroring was broken in the middle of the synchronization process and the slave might be inconsistent.
 - In this case, the administrator must choose whether to use the most updated version, which might be inconsistent, or the last consistent snapshot.

- Reverting the volume to the last consistent snapshot can only be performed by deleting the mirroring, reverting the volume and creating a new mirroring definition.
- In any case, if a last consistent snapshot exists, a most updated snapshot is created, keeping a copy of the information at the time of the role change.

Changing the roles in asynchronous mirroring:

- When applied on the master:
 - Upon successful issuance of the command on the master, the master is reverted to the image recorded on the last replicated snapshot of the mirror, it ceases accepting host requests, and does not accept replication from the other peer as a slave.
- When applied on the slave:
 - A warning is displayed: **Are you sure to change slave to master?**
 - 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 new master establishes asynchronous interval-based sync job process, based on the schedule.
 - Completion of process is recorded in the log.
 - Mirroring state is Standby.
 - Explicit activation of mirroring is required.

Requirements for a successful command completion:

- The command cannot be issued on the master during the Initialization phase.
- The command cannot be issued in the Change Tracking state.
- The activation state is Standby.
- The command can be applied on a volume only if the volume is not part of a mirrored consistency group; if the consistency group is mirrored, the command returns an error and fails.
- The command can be issued on the slave, except during initialization.

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's local peer to become Slave and loose 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's local peer to become Master? The local peer has a last-consistent snapshot

Return codes

- VOLUME_BAD_NAME
Volume name does not exist
- VOLUME_NO_MIRROR
Local volume does not have remote mirroring definitions
- CONS_GROUP_BAD_NAME
Consistency Group name does not exist.
- CONS_GROUP_NO_MIRROR
Local Consistency Group does not have remote mirroring definitions
- MIRROR_IS_INITIAL
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
Volume mirror is part of Consistency Group mirror.
- MIRROR_RETRY_OPERATION
There is an operation in progress on this mirror , please try again your request in a few seconds

Troubleshooting: Please try again the command in a few seconds

- **MIRROR_HAS_NO_SYNCED_SNAPSHOT**
Mirror does not have a synchronized Snapshot.
- **MASTER_CANNOT_BE_DEMOTED**
Master cannot be demoted to Slave role, Peer status mismatch
- **VOLUME_HAS_MULTIPLE_MIRRORS**
Volume has multiple mirrors, operation not allowed or target must be specified
- **TARGET_BAD_NAME**
Target name does not exist
- **VOLUME_TARGET_MISMATCH**
Volume and target do not match
- **CONS_GROUP_BAD_TARGET**
Target name does not match the Consistency Group
- **MIRROR_PART_OF_XMIRROR**
Remote mirror is part of xmirror
- **MIRROR_IS_STANDBY**
Mirror is marked as standby
- **COMMAND_NOT_SUPPORTED_FOR_OLVM_VOLUMES**
This command is not supported for IBM Hyper-Scale Mobility volumes.
- **COMPRESSED_VOLUME_IS_MAPPED**
Mapped compressed volume cannot be a mirror slave
- **COMPRESSION_RECOVERY_IN_PROGRESS**
Operation cannot be completed because another recovery process is in progress

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:

- The system displays the following warning: **Are you sure to change schedule?**
- 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

User Category	Permission
Technicians	Disallowed

Return codes

- **VOLUME_BAD_NAME**
Volume name does not exist
- **CONS_GROUP_BAD_NAME**
Consistency Group name does not exist.
- **VOLUME_NO_MIRROR**
Local volume does not have remote mirroring definitions
- **CONS_GROUP_NO_MIRROR**
Local Consistency Group does not have remote mirroring definitions
- **SCHEDULE_DOES_NOT_EXIST**
Specified Schedule does not exist
- **SYNC_MIRROR_DOES_NOT_USE_SCHEDULE**
Definition of Synchronous Mirror does not require a Schedule object to be specified.
- **VOLUME_BELONGS_TO_MIRRORED_CONS_GROUP**
Volume mirror is part of Consistency Group mirror.
- **INTERVAL_SHOULD_BE_SHORTER_THAN_RPO**
Schedule interval must be shorter than the RPO.
- **TARGET_BAD_NAME**
Target name does not exist
- **VOLUME_HAS_MULTIPLE_MIRRORS**
volume has multiple mirrors, operation not allowed or target must be specified
- **VOLUME_TARGET_MISMATCH**
Volume and target do not match
- **CONS_GROUP_BAD_TARGET**
Target name does not match the Consistency Group
- **MIRROR_IS_STANDBY**
mirror is marked as standby
- **DATA_REDUCTION_TIER_IS_OFFLINE**
Data reduced tier is offline, operation is not allowed

Troubleshooting: Wait for data reduction recovery to start and recover the compression objects

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 ]
```

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]

Name	Type	Description	Mandatory	Default
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]
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_xmirror	Boolean	Marks the mirror as part of xmirror	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 (xmirror) for an existing mirrored volume, use the `part_of_xmirror` 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 relationship 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 relationship 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.

Prior to issuing this command on a consistency group, make sure that the consistency group is empty.

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.

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

- `ASYNC_MIRROR_MISSING_RPO`
Definition of Asynchronous Mirror requires RPO to be specified.
- `ASYNC_MIRROR_REMOTE_RPO_TOO_LONG`
Specified Remote RPO is too long.
- `ASYNC_MIRROR_REMOTE_RPO_TOO_SHORT`
Specified Remote RPO is too short.
- `ASYNC_MIRROR_RPO_TOO_SHORT`
Specified RPO is too short.
- `ASYNC_MIRROR_RPO_TOO_LONG`
Specified RPO is too long.
- `ASYNC_NOT_SUPPORTED_IN_TARGET`
Asynchronous Mirror is not supported by specified Target.
- `BAD_REMOTE_VOLUME_NAME`
Slave volume name does not exist
- `BAD_REMOTE_VOLUME_SIZE`
Master and slave volumes contain a different number of blocks

- **CONS_GROUP_BAD_NAME**
Consistency Group name does not exist.
- **CONS_GROUP_HAS_MIRROR**
Consistency Group has mirroring defined for it.
- **CONS_GROUP_MIRRORING_NOT_SUPPORTED_IN_TARGET**
Consistency Group mirroring is not supported by target machine.
- **INTERVAL_SHOULD_BE_SHORTER_THAN_RPO**
Schedule interval must be shorter than the RPO.
- **MAX_MIRRORS_REACHED**
Maximum number of mirrors already defined
- **NOT_ENOUGH_SPACE_ON_REMOTE_MACHINE**
Not enough free space to set requested size of slave volume
- **VOLUME_BAD_NAME**
Volume name does not exist
- **VOLUME_IS_MASTER**
Local volume is already defined as a master volume
- **VOLUME_IS_SLAVE**
Volume is defined as a slave volume
- **REMOTE_VOLUME_EXISTS**
Slave volume name already exists and cannot be created
- **REMOTE_MAX_VOLUMES_REACHED**
Maximum number of volumes already defined on remote machine
- **REMOTE_MAX_MIRRORS_REACHED**
Maximum number of mirrors already defined on remote machine
- **VOLUME_BAD_PREFIX**
Volume name has a reserved prefix
- **REMOTE_POOL_DOES_NOT_EXIST**
Pool does not exist on remote machine
- **REMOTE_POOL_NOT_SPECIFIED**
A Pool on remote machine must be specified when a slave volume is to be created
- **REMOTE_TARGET_NOT_CONNECTED**
There is currently no connection from the target system
- **VOLUME_IS_SNAPSHOT**
Operation is not permitted on snapshots
- **REMOTE_VOLUME_IS_SNAPSHOT**
Slave volume is a snapshot
- **TARGET_BAD_NAME**
Target name does not exist
- **TARGET_BAD_TYPE**
Target machine is not XIV machine
- **TARGET_NO_ACCESS**
No access permissions to slave machine
- **TARGET_NOT_CONNECTED**
There is currently no connection to the target system
- **REMOTE_VOLUME_LOCKED**
Slave volume is locked
- **TIMEOUT**
Remote operation did not complete in time
- **VOLUME_HAS_MIRRORING_SNAPSHOTS**
Volume has snapshots created by previous mirroring process.
- **SLAVE_VOLUME_NOT_FORMATTED**
Slave volume is not formatted
- **TARGET_DOES_NOT_ACCEPT_XIV_COMMANDS**
Target system does not accept XIV management commands
- **SYNC_MIRROR_HAS_NO_RPO**
Synchronous Mirror does not have an RPO.

- REMOTE_CONS_GROUP_IS_MIRRORED
Remote Consistency Group has mirroring defined for it.
- REMOTE_SCHEDULE_DOES_NOT_EXIST
Specified Schedule does not exist on remote machine
- SCHEDULE_DOES_NOT_EXIST
Specified Schedule does not exist
- REMOTE_CONS_GROUP_BAD_NAME
Remote Consistency Group name does not exist.
- REMOTE_VOLUME_IS_MASTER
Volume on remote machine is currently defined as Master
- REMOTE_VOLUME_IS_SLAVE
Slave volume is already defined as a slave volume
- REMOTE_MAX_MIRROR_CAPACITY_REACHED
Maximum capacity for mirrored volumes already defined on remote machine
- MIRROR_RETRY_OPERATION
There is an operation in progress on this mirror , please try again your request in a few seconds

Troubleshooting: Please try again the command in a few seconds

- MIRRORING_INCOMPATIBLE_TARGET_VERSION
Mirroring is not supported between the system versions of the specified peers.
- NO_OFFLINE_INIT_TYPE_WITH_SLAVE_CREATION
New Volume will be created as slave. Offline init meaningless.
- ASYNC_WITH_OFFLINE_INIT_NOT_SUPPORTED_IN_TARGET
Asynchronous Mirror with offline initialization option is not supported by the specified Target.
- VOLUME_SIZE_ABOVE_LIMIT
Volume size specified is above limit
- REMOTE_VOLUME_SIZE_ABOVE_LIMIT
Volume size specified is above limit of remote machine
- INVALID_SLICE_OFFSET
Slice offset is illegal
- SLAVE_VOLUME_NOT_SAME_TYPE
Slave volume is not of the same type (compressed or uncompressed) as the master
- COMPRESSION_NOT_SUPPORTED_IN_TARGET
Compressed volume is not supported by the specified Target.
- OFFLINE_INIT_NOT_SUPPORTED_IN_COMPRESSED_VOLUME
Mirror with offline initialization option is not supported with compressed volumes.
- VOLUME_HAS_TRANSFORM
Operation is not supported for volumes which are compressed/decompressed
- MAX_COMPRESSED_VOLUMES_REACHED
Maximum number of compressed volumes already defined
- 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
System is in the process of changing encryption activation state
- MIRROR_OF_SAME_TYPE_EXISTS_ON_VOLUME
A mirror of the same type already defined on this volume
- XMIRROR_IS_NOT_SUPPORTED_FOR_CONS_GROUPS
A CG Cannot be define part of xmirror
- MIRROR_EXISTS_ON_TARGET
Volume already has a mirror on this target
- REMOTE_VOLUME_IS_MIRROR_MASTER
Volume is a mirror master. Can't be slave!
- XMIRROR_MAX_NUM_OF_MIRRORS_REACHED
Failed to create mirror, max number of mirrors exceeded
- REMOTE_VOLUME_TWO_SYNC_MIRRORS_NOT_ALLOWED

Two SYNC mirrors detected on 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 and its target is this system so you can't create a mirror with this target here.
- **VOLUME_BELONGS_TO_MIRRORED_CONS_GROUP**
Volume mirror is part of Consistency Group mirror.
- **DOMAIN_MAX_MIRRORS_REACHED**
The domain exceeds the maximum allowed number of mirrors.
- **REMOTE_DOMAIN_MAX_VOLUMES_REACHED**
Maximum number of volumes already defined on remote machine domain
- **REMOTE_DOMAIN_HAS_NO_ACCESS_TO_TARGET**
Slave machine domain has no access to target
- **REMOTE_DOMAIN_HAS_NO_ACCESS_TO_SCHEDULE**
Slave machine domain has no access to schedule
- **DOMAIN_HAS_NO_ACCESS_TO_TARGET**
Domain has no access to target.
- **REMOTE_DOMAIN_MAX_MIRRORS_REACHED**
Maximum number of mirrors already defined on remote machine domain
- **DOMAIN_MAX_VOLUMES_REACHED**
The domain exceeds the maximum allowed number of volumes.
- **REMOTE_VOLUME_HAS_DATA_MIGRATION**
Data Migration is defined for slave volume
- **REMOTE_VOLUME_MASTER_ASYNC_MIRROR_DETECTED**
An ASYNC master mirror was detected on the remote volume. Operation not allowed.
- **MAX_XMIRRORS_REACHED**
The number of Xmirrors objects exceeded limit
- **XMIRROR_MIRRORING_INCOMPATIBLE_TARGET_VERSION**
Xmirror Mirroring is not supported between the system versions of the specified peers.
- **REMOTE_COMPRESSION_REQUIRES_THIN_PROVISIONED_POOL**
Remote compressed volumes creation require thin provisioned pools
- **REMOTE_COMPRESSED_CAPACITY_LIMIT_REACHED**
Total capacity of compressed volumes in the remote exceeds the maximum limit
- **REMOTE_COMPRESSION_IS_NOT_ACTIVE**
Compression is not enabled in the remote side
- **REMOTE_MAX_COMPRESSED_VOLUMES_REACHED**
Maximum number of compressed volumes already defined in the remote
- **REMOTE_COMPRESSED_VOLUME_IS_MAPPED**
Mapped compressed volume in the remote cannot be a mirror slave
- **COMPRESSION_VOLUME_IS_OFFLINE**
Operation cannot be completed because the requested compressed volume is offline
- **REMOTE_VOLUME_IS_OFFLINE**
Operation cannot be completed because slave compressed volume is offline
- **REMOTE_COMPRESSED_VOLUME_TOO_SMALL**
Volume size is below the minimum size allowed for compressed volume in the remote system

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

Name	Type	Description	Mandatory	Default
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
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- VOLUME_BAD_NAME
Volume name does not exist
- VOLUME_NO_MIRROR
Local volume does not have remote mirroring definitions
- CONS_GROUP_BAD_NAME
Consistency Group name does not exist.
- CONS_GROUP_NO_MIRROR
Local Consistency Group does not have remote mirroring definitions
- LOCAL_PEER_IS_NOT_MASTER
Local peer is not the master
- SYNC_ALREADY_INACTIVE
Synchronization is already inactive
- VOLUME_BELONGS_TO_MIRRORED_CONS_GROUP
Volume mirror is part of Consistency Group mirror.
- MIRROR_RETRY_OPERATION
There is an operation in progress on this mirror , please try again your request in a few seconds
Troubleshooting: Please try again the command in a few seconds
- TARGET_BAD_NAME
Target name does not exist
- VOLUME_HAS_MULTIPLE_MIRRORS
Volume has multiple mirrors, operation not allowed or target must be specified
- VOLUME_TARGET_MISMATCH
Volume and target do not match
- CONS_GROUP_BAD_TARGET
Target name does not match the Consistency Group
- REMOTE_MIRROR_IS_STANDBY
Remote mirror is marked as standby
- MIRROR_IS_STANDBY

mirror is marked as standby

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]
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 relationships of the CG and of all volumes in the CG?

Return codes

- VOLUME_BAD_NAME
Volume name does not exist
- VOLUME_NO_MIRROR
Local volume does not have remote mirroring definitions
- CONS_GROUP_BAD_NAME
Consistency Group name does not exist.

- **CONS_GROUP_NO_MIRROR**
Local Consistency Group does not have remote mirroring definitions
- **LOCAL_PEER_IS_NOT_MASTER**
Local peer is not the master
- **MIRROR_IS_ACTIVE**
Remote mirroring is currently active
- **FORCE_DELETE_NOT_ALLOWED_ON_MASTER**
Only slave mirrors need to be forced to be deleted
- **VOLUME_BELONGS_TO_MIRRORED_CONS_GROUP**
Volume mirror is part of Consistency Group mirror.
- **MIRROR_RETRY_OPERATION**
There is an operation in progress on this mirror , please try again your request in a few seconds
- **Troubleshooting:** Please try again the command in a few seconds
- **MIRROR_IS_NOT_INITIALIZING**
Operation is permitted only during the Initialization phase.
- **TARGET_BAD_NAME**
Target name does not exist
- **VOLUME_HAS_MULTIPLE_MIRRORS**
volume has multiple mirrors, operation not allowed or target must be specified
- **VOLUME_TARGET_MISMATCH**
Volume and target do not match
- **MIRROR_ASSOCIATED_WITH_XMIRROR**
This mirror is associated with a defined xmirror, operation not allowed
- **CONS_GROUP_BAD_TARGET**
Target name does not match the Consistency Group
- **REMOTE_MIRROR_IS_STANDBY**
Remote mirror is marked as standby
- **MIRROR_IS_STANDBY**
Mirror is marked as standby

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 available 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** or **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 (in MB)
- Operational: **Yes** or **No**
- Sync Progress (in %)
- Mirror Error: specifies the reason for mirroring deactivation: **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):

- **INACTIVE_USER - No_Error**
- **INACTIVE_SECONDARY_LOCKED - Secondary_Pool_Exhausted**
- **INACTIVE_POOL_EXHAUSTED - Master_Pool_Exhausted**
- **INACTIVE_VOL_SIZE_MISMATCH - Remote_And_Local_Volume_Size_Mismatch**
- **INACTIVE_CONS_GROUP_MEMBERSHIP_MISMATCH - Cons_Group_Membership_Mismatch**
- **INACTIVE_POSSIBLE_VOL_SIZE_MISMATCH - Possible_Remote_And_Local_Volume_Size_Mismatch**
- **INACTIVE_POSSIBLE_CONS_GROUP_MEMBERSHIP_MISMATCH - Possible_Cons_Group_Membership_Mismatch**
- **INACTIVE_THIN_PROVISIONING - No_Thin_Provisioning_Resources**
- **INACTIVE_PEER_STATUS_MISMATCH - Peer_Status_Mismatch**
- **INACTIVE_UPGRADE - Temporarily_Deactivated_For_Upgrade**

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
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
application_consistent	App Consistency	N/A	N/A
validate	Validation	N/A	N/A
is_standby	Standby	N/A	8

Output:

```
<command id=" home_markdown_jenkins_workspace_Transform_in_STJTAG_Gen3_XCLI_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_STJTAG_Gen3_XCLI_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"/>
      </mirror>
    </command>
  </administrator>
</command>
```

```

    <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 (MB)

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_STJTAG_Gen3_XCLI_r_command_mirroring_mirror_statistics_get_143">
  <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
Volume name does not exist
- CONS_GROUP_NO_MIRROR
Local Consistency Group does not have remote mirroring definitions
- MIRROR_HAS_NO_STATISTICS
Job statistics were not collected for this Mirror.
- LOCAL_IS_SLAVE
Local mirror peer is not the master
- VOLUME_BELONGS_TO_MIRRORED_CONS_GROUP
Volume mirror is part of Consistency Group mirror.
- VOLUME_NO_MIRROR
Local volume does not have remote mirroring definitions
- CONS_GROUP_BAD_NAME
Consistency Group name does not exist.
- SYNC_MIRROR_HAS_NO_STATISTICS
Job statistics do not exist for Synchronous Mirror.
- TARGET_BAD_NAME
Target name does not exist
- VOLUME_HAS_MULTIPLE_MIRRORS
volume has multiple mirrors, operation not allowed or target must be specified
- VOLUME_TARGET_MISMATCH
Volume and target do not match
- CONS_GROUP_BAD_TARGET
Target name does not match the Consistency Group
- MIRROR_IS_STANDBY
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

Return codes

- **VOLUME_BAD_NAME**
Volume name does not exist
- **VOLUME_NO_MIRROR**
Local volume does not have remote mirroring definitions
- **CONS_GROUP_BAD_NAME**
Consistency Group name does not exist.
- **CONS_GROUP_NO_MIRROR**
Local Consistency Group does not have remote mirroring definitions
- **LOCAL_PEER_IS_NOT_MASTER**
Local peer is not the master
- **MIRROR_IS_NOT_SYNCHRONIZED**
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**
Volume mirror is part of Consistency Group mirror.
- **MIRROR_HAS_SYNC_JOB**
Operation is not permitted on a mirror with active sync jobs
- **MIRROR_RETRY_OPERATION**
There is an operation in progress on this mirror , please try again your request in a few seconds

Troubleshooting: Please try again the command in a few seconds
- **MIRROR_MASTER_DIFFERS_FROM_SLAVE**
Mirror master was written to after the last replicated snapshot was taken
- **REMOTE_MIRROR_IS_NOT_ACTIVE**
Remote mirroring is not active
- **TARGET_BAD_NAME**
Target name does not exist
- **VOLUME_HAS_MULTIPLE_MIRRORS**
volume has multiple mirrors, operation not allowed or target must be specified
- **VOLUME_TARGET_MISMATCH**
Volume and target do not match
- **CONS_GROUP_BAD_TARGET**
Target name does not match the Consistency Group
- **REMOTE_MIRROR_IS_STANDBY**
remote mirror is marked as standby
- **MIRROR_IS_STANDBY**
mirror is marked as standby
- **COMMAND_NOT_SUPPORTED_FOR_OLVM_VOLUMES**
This command is not supported for IBM Hyper-Scale Mobility volumes.
- **COMPRESSED_VOLUME_IS_MAPPED**

Mapped compressed volume cannot be a mirror slave

- **COMPRESSION_RECOVERY_IN_PROGRESS**
Operation cannot be completed because another recovery process is in progress

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

Example:

rpo_thresholds_get

Output:

```
Increase Percentage  Increase Absolute
-----
100                  3600
```

Field ID	Field output	Default position
increase_percentage	Increase Percentage	1
increase_absolute	Increase Absolute	2

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
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 executed 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**

Values should be in [1,10000] range

- **INVALID_RPO_THRESHOLD_ABSOLUTE**
Values should be in [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:

- The schedule must be one of the following values: 00:00:30, 00:01, 00:02, 00:05, 00:10, 00:15, 00:30, 01:00, 02:00, 03:00, 06:00, 08:00, 12:00.
- 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_create interval=00:01 schedule=1min domain=* -y
```

Output:

```
Command executed 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_UPDATE_THE_SCHEDULE**
Are you sure you want to update this schedule? This change will effect all mirrors using that schedule.

Return codes

- **SCHEDULE_DOES_NOT_EXIST**
Specified Schedule does not exist
- **BAD_SCHEDULE_TIME_FORMAT**
Time format for Schedule is HH:MM[:SS]
- **ILLEGAL_INTERVAL**
Specified interval value is not supported.
- **SCHEDULE_CAN_NOT_BE_UPDATED**
Specified Schedule cannot be updated
- **INTERVAL_SCHEDULE_REQUIRES_ONLY_ONE_INTERVAL**
Multiple times should not be defined for Interval Schedule
- **SCHEDULE_EXCLUDE_TIMES_NOT_REQUIRED**
Exclusion period may defined only is exclude_time is set

- ZERO_LENGTH_EXCLUSION_PERIOD
Exclusion period start time must be different than its end time
- DOMAIN_SCHEDULE_IN_USE
Cannot move the schedule to other domain since it is in use.
- DOMAIN_DOESNT_EXIST
Domain does not exist.

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: 00:00:30, 00:01, 00:02, 00:05, 00:10, 00:15, 00:30, 01:00, 02:00, 03:00, 06:00, 08:00, 12:00.
- 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=lmin domain=*
```

Output:

```
Command executed 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
Schedule name exists
- BAD_SCHEDULE_TIME_FORMAT
Time format for Schedule is HH:MM[:SS]

- **MAX_SYNC_SCHEDULES_REACHED**
Maximal number of schedule objects has been reached
- **ILLEGAL_INTERVAL**
Specified interval value is not supported.
- **INTERVAL_SCHEDULE_REQUIRES_ONLY_ONE_INTERVAL**
Multiple times should not be defined for Interval Schedule
- **ZERO_LENGTH_EXCLUSION_PERIOD**
Exclusion period start time must be different than its end time
- **SCHEDULE_EXCLUDE_TIMES_NOT_REQUIRED**
Exclusion period may be defined only if exclude_time is set
- **ONLY_INTERVAL_SCHEDULE_MAY_HAVE_EXCLUSIONS**
Exclusion period may be defined only for Interval Schedule
- **DOMAIN_DOESNT_EXIST**
Domain does not exist.

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**
Specified Schedule does not exist
- **MAX_VOLUMES_REACHED**
Maximum number of volumes already defined
- **SCHEDULE_IS_NOT_MANUAL**
Specified Schedule does not allow 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 executed 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
Specified Schedule is currently assigned to a mirror
- SCHEDULE_CAN_NOT_BE_DELETED
Specified Schedule cannot be deleted
- SCHEDULE_DOES_NOT_EXIST
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
Specified Schedule does not exist
- SCHEDULE_NAME_EXISTS
New Schedule name already exists
- SCHEDULE_CAN_NOT_BE_UPDATED
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: initialization, pending, running, complete
- Type: interval-initiated, Snapshot Mirror, initialization, initializing validate
- 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

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 the 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

User Category	Permission
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- VOLUME_BAD_NAME
Volume name does not exist
- VOLUME_NO_DM
Local volume does not have Data Migration definitions
- TARGET_NOT_CONNECTED
There is currently no connection to the target system
- REMOTE_VOLUME_NO_LUN
Remote volume's LUN is unavailable
- REMOTE_VOLUME_NO_READ_ACCESS
Remote volume cannot be read
- REMOTE_VOLUME_NO_WRITE_ACCESS
Remote volume is write protected
- BAD_REMOTE_VOLUME_SIZE
Master and slave volumes contain a different number of blocks

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, data migration can be deleted if it is done
- ARE_YOU_SURE_YOU_WANT_TO_DEACTIVATE_SOURCE_UPDATING_DATA_MIGRATION
Deactivation may cause loss of access to hosts, and will stop all applications. Data migration can be deleted if it is done

Return codes

- VOLUME_BAD_NAME
Volume name does not exist
- VOLUME_NO_DM
Local volume does not have Data Migration definitions

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 relationship 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

Return codes

- **VOLUME_BAD_NAME**
Volume name does not exist
- **TARGET_BAD_NAME**
Target name does not exist
- **VOLUME_IS_SNAPSHOT**
Operation is not permitted on snapshots
- **VOLUME_HAS_MIRROR**
Mirror is defined for this volume
- **VOLUME_BELONGS_TO_CG**
Volume belongs to a Consistency Group
- **VOLUME_HAS_DATA_MIGRATION**
Data Migration is defined for this volume
- **VOLUME_HAS_SNAPSHOTS**
Volume has snapshots
- **VOLUME_NOT_FORMATTED**
Local volume is not formatted
- **VOLUME_EXISTS**
Volume name already exists
- **POOL_DOES_NOT_EXIST**
Storage Pool does not exist

- **VOLUME_BAD_PREFIX**
Volume name has a reserved prefix
 - **NOT_ENOUGH_SPACE**
No space to allocate for volume's current usage
 - **NOT_ENOUGH_HARD_SPACE**
No space to allocate for volume's current usage
 - **MAX_VOLUMES_REACHED**
Maximum number of volumes already defined
 - **REMOTE_VOLUME_NO_LUN**
Remote volume's LUN is unavailable
 - **TARGET_NOT_CONNECTED**
There is currently no connection to the target system
 - **VOLUME_CANNOT_HAVE_ZERO_SIZE**
Volume size cannot be zero
 - **ILLEGAL_LUN**
LUN is out of range
 - **NO_ONLINE_MIGRATION_WITHOUT_SOURCE_UPDATING**
Data Migration without automatic migration must be defined as source-updating
 - **MIGRATION_ALREADY_DEFINED_FOR_LUN**
Data Migration is already defined from lun *LUN* of target '*Target*'
 - **VOLUME_SIZE_ABOVE_LIMIT**
Volume size specified is above limit
 - **INVALID_SLICE_OFFSET**
Slice offset is illegal
 - **DM_COMPRESSED_NOT_SUPPORTED**
Compressed volume can not be target for Data Migration'
 - **VOLUME_HAS_TRANSFORM**
Operation is not supported for volumes which are compressed/decompressed
 - **ENCRYPTION_IN_PROGRESS**
System is in the process of changing encryption activation state
 - **DOMAIN_MAX_VOLUMES_REACHED**
The domain exceeds the maximum allowed number of volumes.
 - **MAX_DMS_REACHED**
Maximum number of remote volumes (mirror/migration) is already defined
- Troubleshooting:** Delete unnecessary Data Migration objects
- **DOMAIN_MAX_DMS_REACHED**
The domain exceeds the maximum allowed number of data migrations.

Deleting the data migration process

Use the **dm_delete** command to delete the data migration process.

dm_delete vol=VolName

Parameters

Name	Type	Description	Mandatory
vol	Object name	The name of the volume whose data migration process is to be deleted.	Y

This command deletes the data migration configuration and stops the data migration process.

This command can only be executed if the data migration has reached the state of synchronization.

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
Volume name does not exist
- VOLUME_NO_DM
Local volume does not have Data Migration definitions
- DM_IS_NOT_SYNCHRONIZED
Data Migration process has not been completed

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

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
Volume name does not exist
- VOLUME_NO_DM
Local volume does not have Data Migration definitions
- TARGET_NOT_CONNECTED
There is currently no connection to the target system
- REMOTE_VOLUME_NO_LUN
Remote volume's LUN is unavailable
- REMOTE_VOLUME_NO_READ_ACCESS
Remote volume cannot be read
- REMOTE_VOLUME_NO_WRITE_ACCESS
Remote volume is write protected
- BAD_REMOTE_VOLUME_SIZE
Master and slave volumes contain a different number of blocks

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 relationship.
- [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

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**
Volume size is very large. It may not be possible to mirror this volume to older versions of XIV. Are you sure?

Return codes

- **VOLUME_NOT_APPLICABLE_FOR_OLVM**
The volume is not applicable for IBM Hyper-Scale Mobility.
- **VOLUME_BAD_NAME**
Volume name does not exist
- **VOLUME_IS_SNAPSHOT**
Operation is not permitted on snapshots
- **TARGET_BAD_NAME**
Target name does not exist
- **VOLUME_IS_MASTER**
Local volume is already defined as a master volume
- **TARGET_BAD_TYPE**
Target machine is not XIV machine
- **TARGET_NO_ACCESS**
No access permissions to slave 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**
Maximum number of migrations already defined
- **REMOTE_MAX_MIGRATIONS_REACHED**
Maximum number of migrations already defined on remote machine
- **REMOTE_POOL_DOES_NOT_EXIST**
Pool does not exist on remote machine
- **BAD_REMOTE_VOLUME_SIZE**
Master and slave volumes contain a different number of blocks
- **NOT_ENOUGH_SPACE_ON_REMOTE_MACHINE**
Not enough free space to set requested size of slave volume
- **REMOTE_VOLUME_EXISTS**
Slave volume name already exists and cannot be created
- **REMOTE_VOLUME_IS_MASTER**
Volume on remote machine is currently defined as Master
- **REMOTE_VOLUME_IS_SLAVE**
Slave volume is already defined as a slave volume
- **REMOTE_MAX_VOLUMES_REACHED**
Maximum number of volumes already defined on remote machine

- **TIMEOUT**
Remote operation did not complete in time
- **VOLUME_BAD_PREFIX**
Volume name has a reserved prefix
- **REMOTE_VOLUME_HAS_DATA_MIGRATION**
Data Migration is defined for slave volume
- **REMOTE_VOLUME_LOCKED**
Slave volume is locked
- **VOLUME_HAS_MIRRORING_SNAPSHOTS**
Volume has snapshots created by previous mirroring process.
- **REMOTE_MAX_MIRROR_CAPACITY_REACHED**
Maximum capacity for mirrored volumes already defined on remote machine
- **TARGET_DOES_NOT_ACCEPT_XIV_COMMANDS**
Target system does not accept XIV management commands
- **MAX_VOLUMES_REACHED**
Maximum number of volumes already defined
- **VOLUME_LOCKED**
Volume is locked
- **NO_ASYNC_IN_THIN_PROVISIONED_POOL**
Thin provisioned Pool cannot contain Volumes with Asynchronous Mirroring
- **BAD_REMOTE_VOLUME_NAME**
Slave volume name does not exist
- **REMOTE_VOLUME_SIZE_ABOVE_LIMIT**
Volume size specified is above limit of remote machine
- **MIRROR_RETRY_OPERATION**
There is an operation in progress on this mirror , please try again your request in a few seconds

Troubleshooting: Please try again 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**
Local volume is not formatted
- **MIRRORING_INCOMPATIBLE_TARGET_VERSION**
Mirroring is not supported between the system versions of the specified peers.
- **NOT_ENOUGH_SPACE**
No space to allocate volume
- **VOLUME_SIZE_ABOVE_LIMIT**
Volume size specified is above limit
- **INVALID_SLICE_OFFSET**
Slice offset is illegal
- **ILLEGAL_VOLUME_SIZE**
Illegal volume size
- **VOLUME_IS_OLVM_PROXY**
The volume is in an IBM Hyper-Scale Mobility Proxy phase.
- **VOLUME_IS_SLAVE**
Volume is defined as a slave volume
- **REMOTE_VOLUME_IS_SNAPSHOT**
Slave volume is a snapshot
- **VOLUME_EXISTS**
Volume name already exists
- **SLAVE_VOLUME_NOT_FORMATTED**
Slave volume is not formatted
- **VOLUME_BELONGS_TO_CG**
Volume belongs to a Consistency Group
- **VOLUME_HAS_DATA_MIGRATION**
Data Migration is defined for this volume

- **MAX_MIRRORS_REACHED**
Maximum number of mirrors already defined
- **VOLUME_CANNOT_HAVE_ZERO_SIZE**
Volume size cannot be zero
- **ASYNC_MIRROR_REMOTE_RPO_TOO_SHORT**
Specified Remote RPO is too short.
- **POOL_DOES_NOT_EXIST**
Storage Pool does not exist
- **REMOTE_VOLUME_NOT_APPLICABLE_FOR_OLVM**
The remote volume is not applicable for IBM Hyper-Scale Mobility.
- **REMOTE_SCHEDULE_DOES_NOT_EXIST**
Specified Schedule does not exist on remote machine
- **OLVM_DOES_NOT_SUPPORT_ISCSI_TARGET**
IBM Hyper-Scale Mobility does not support ISCSI targets.
- **ASYNC_MIRROR_REMOTE_RPO_TOO_LONG**
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**
System is in the process of changing encryption activation state
- **MAX_OLVM_REACHED**
Maximum number of IBM Hyper-Scale Mobility Relations already defined
- **DOMAIN_MAX_MIRRORS_REACHED**
The domain exceeds the maximum allowed number of mirrors.
- **REMOTE_DOMAIN_MAX_MIGRATIONS_REACHED**
Maximum number of migrations already defined on remote machine domain
- **DOMAIN_HAS_NO_ACCESS_TO_TARGET**
Domain has no access to target.
- **REMOTE_DOMAIN_HAS_NO_ACCESS_TO_TARGET**
Slave machine domain has no access to target
- **DOMAIN_MAX_VOLUMES_REACHED**
The domain exceeds the maximum allowed number of volumes.
- **REMOTE_DOMAIN_MAX_VOLUMES_REACHED**
Maximum number of volumes already defined on remote machine domain
- **OLVM_BLACKLIST_FULL**
Cannot create new IBM Hyper-Scale Mobility Relations. Too many volume serials are blacklisted.
- **XMIRROR_MAX_NUM_OF_MIRRORS_REACHED**
Failed to create mirror, max number of mirrors exceeded
- **REMOTE_VOLUME_TWO_SYNC_MIRRORS_NOT_ALLOWED**
Two SYNC mirrors detected on remote volume. This is not allowed.
- **REMOTE_VOLUME_IS_MIRROR_MASTER**
Volume is a mirror master. Can't be slave!
- **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 and its target is this system so you can't create a mirror with this target here.
- **REMOTE_VOLUME_MASTER_ASYNC_MIRROR_DETECTED**
An ASYNC master mirror was detected on the remote volume. Operation not allowed.
- **XMIRROR_MIRRORING_INCOMPATIBLE_TARGET_VERSION**
Xmirror Mirroring is not supported between the system versions of the specified peers.
- **VOLUME_HAS_MULTIPLE_MIRRORS**
volume has multiple mirrors, operation not allowed or target must be specified
- **SLAVE_VOLUME_NOT_SAME_TYPE**
Slave volume is not of the same type (compressed or uncompressed) as the master
- **VOLUME_HAS_TRANSFORM**
Operation is not supported for volumes which are compressed/decompressed
- **REMOTE_MAX_COMPRESSED_VOLUMES_REACHED**

Maximum number of compressed volumes already defined in the remote

- **REMOTE_COMPRESSED_CAPACITY_LIMIT_REACHED**
Total capacity of compressed volumes in the remote exceeds the maximum limit
- **REMOTE_COMPRESSION_REQUIRES_THIN_PROVISIONED_POOL**
Remote compressed volumes creation require thin provisioned pools
- **COMPRESSION_NOT_SUPPORTED_IN_TARGET**
Compressed volume is not supported by the specified Target.
- **REMOTE_COMPRESSION_IS_NOT_ACTIVE**
Compression is not enabled in the remote side
- **COMPRESSION_VOLUME_IS_OFFLINE**
Operation cannot be completed because the requested compressed volume is offline
- **REMOTE_VOLUME_IS_OFFLINE**
Operation cannot be completed because slave compressed volume is offline
- **REMOTE_COMPRESSED_VOLUME_TOO_SMALL**
Volume size is below the minimum size allowed for compressed volume in the remote system

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 relationship.

```
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**
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
- **MIRROR_CONFIGURATION_ERROR**
Mirror local configuration does not match remote configuration
- **MIRROR_CONS_GROUP_MEMBERSHIP_MISMATCH**
Mirrored CG contains different volumes on Master and Slave. This problem occurs whenever the `cg_add_vol` or `cg_remove_vol` commands were previously issued and the Master did not receive an acknowledgment from the Slave until the command timed out, or any other unexpected failure.
- **MIRROR_RETRY_OPERATION**
There is an operation in progress on this mirror , please try again your request in a few seconds

Troubleshooting: Please try again 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**
Slave volume and Master Volume sizes are different
- **REMOTE_VOLUME_IS_MASTER**
Volume on remote machine is currently defined as Master
- **REMOTE_MAX_VOLUMES_REACHED**

Maximum number of volumes already defined on remote machine

- REMOTE_MAX_COMPRESSED_VOLUMES_REACHED
Maximum number of compressed volumes already defined in the remote
- VOLUME_NOT_DEFINED_FOR_OLVM
The volume does not have IBM Hyper-Scale Mobility definitions
- OLVM_ALREADY_ACTIVE
The IBM Hyper-Scale Mobility relation is already active.
- MAX_VOLUMES_REACHED
Maximum number of volumes already defined
- REMOTE_MIRROR_IS_STANDBY
remote mirror is marked as standby
- VOLUME_HAS_MULTIPLE_MIRRORS
volume has multiple mirrors, operation not allowed or target must be specified
- REMOTE_DOMAIN_MAX_VOLUMES_REACHED
Maximum number of volumes already defined on remote machine domain
- COMPRESSION_VOLUME_IS_OFFLINE
Operation cannot be completed because the requested compressed volume is offline

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
Volume name does not exist
- REMOTE_VOLUME_IS_MASTER
Volume on remote machine is currently defined as Master
- 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 relation is already inactive.
- REMOTE_MIRROR_IS_STANDBY

remote mirror is marked as standby

- **VOLUME_HAS_MULTIPLE_MIRRORS**
volume has multiple mirrors, operation not allowed or target must be specified

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 relationship on the source.	N	No
force_abort_on_destination	Boolean	Determine whether to delete an IBM Hyper-Scale Mobility relationship 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 relationship. 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
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 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 IBM Hyper-Scale Mobility relationship?

Return codes

- **VOLUME_BAD_NAME**
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 relation is active.
- **FORCE_ABORT_NOT_ALLOWED**
A forced IBM Hyper-Scale Mobility abort is not allowed.
- **VOLUME_HAS_MULTIPLE_MIRRORS**
volume has multiple mirrors, operation not allowed or target must be specified

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 a Proxy state? Source volume and all volume snapshots will be deleted.

Return codes

- VOLUME_BAD_NAME
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
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 Proxy state started
- HOST_BAD_NAME
Host name does not exist
- ISCSI_HOST_ILLEGAL_PORT_NAME
Port name for iSCSI Host is illegal

Troubleshooting: Port names for iSCSI Hosts must contain only printable characters.
- MAX_PORTS_REACHED
Maximum number of ports already defined in the system
- HOST_PORT_EXISTS
Host with this port ID already defined
- REMOTE_MAX_VIRTUAL_HOSTS_REACHED
Maximum number of remote virtual hosts already defined
- OLVM_RETRY_OPERATION
There is an operation in progress on this olvm , please try again your request in a few seconds

Troubleshooting: Please try again the command in a few seconds
- VOLUME_HAS_MULTIPLE_MIRRORS
volume has multiple mirrors, operation not allowed or target must be specified

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 relationship 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
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 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 IBM Hyper-Scale Mobility relationship?

Return codes

- VOLUME_BAD_NAME
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
Volume that is mapped to a host cannot be deleted
- VOLUME_IS_BOUND
Volume is bound to a ALU
Troubleshooting: Unbound the volume from the ALU
- VOLUME_HAS_MULTIPLE_MIRRORS
volume has multiple mirrors, operation not allowed or target must be specified

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 relationships 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
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

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.
- [Generating a custom event](#)
Use the `mm_event` command to generate a maintenance module event.
- [Retrieving a maintenance module's heartbeat](#)
Use the `mm_heartbeat` command to retrieve the heartbeat of a maintenance module.
- [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 prioritize 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 user control event](#)
Use `xmpns_user_config_set` command to generate an `XMPNS_USER_CONTROL` event.
- [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	Must be specified for an 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 executed 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 executed 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> | msgws=<MSGW1[,MSGW2]...|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
smmsgws	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.

- *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 `smtpgw_prioritize` command 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@yourcompany.com
```

Output:

```
Command executed successfully.
```

Example:

```
dest_define dest=monitoringserver type=SNMP
snmp_manager=10.170.68.111
```

Output:

```
Command executed 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
Maximum number of destinations already defined
- DEST_NAME_ALREADY_EXISTS

Destination name already exists

- **DEST_NAME_IS_DESTGROUP_NAME**
Destination name already exists as a destination group name
- **EMAIL_NOT_ALLOWED_FOR_DEST_TYPE**
Destination cannot have an email address
- **GATEWAY_NAME_APPEARS_TWICE**
Gateway name appears twice in the list
- **GATEWAY_NAME_DOES_NOT_EXIST**
Gateway name does not exist
- **MSGWS_NOT_ALLOWED_FOR_DEST_TYPE**
Destination cannot have SMS gateways
- **SMTPGWS_NOT_ALLOWED_FOR_DEST_TYPE**
Destination cannot have SMTP gateways
- **SNMP_MANAGER_MUST_BE_SPECIFIED_FOR_DEST_TYPE**
Destination must have an SNMP manager
- **SNMP_MANAGER_NOT_ALLOWED_FOR_DEST_TYPE**
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**
Destination cannot have an https address
- **PROXY_ADDRESS_NOT_ALLOWED_FOR_DEST_TYPE**
Destination cannot have a proxy address
- **SNMP_DESTS_CANNOT_REFER_TO_USERS**
SNMP destinations cannot refer to users
- **HTTPS_DESTS_CANNOT_REFER_TO_USERS**
HTTPS destinations cannot refer to users
- **NO_SMTP_GATEWAYS_ARE_DEFINED**
An email destination cannot be defined if no SMTP gateways are defined
- **USER_EMAIL_ADDRESS_IS_NOT_DEFINED**
User's email address is not defined
- **USER_PHONE_NUMBER_IS_NOT_DEFINED**
User's phone number is not defined
- **USER_NAME_DOES_NOT_EXIST**
User name does not exist
- **INTERNAL_DESTS_CANNOT_REFER_TO_USERS**
Internal destinations cannot refer to users
- **DAY_APPEARS_TWICE**
The day 'Day' appears twice on the list.

Troubleshooting: Each day must appear at most once.

- **HTTPS_ADDRESS_MUST_BE_SPECIFIED_FOR_DEST_TYPE**
Destination must have an https address
- **DEST_TYPE_NOT_SUPPORTED**
This type of destination is not supported

Troubleshooting: Contact support

- **USER_IS_NOT_IN_DESTINATION_DOMAINS**
User must be included in the destination domains.
- **DOMAIN_DOESNT_EXIST**
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 executed 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
Are you sure you want to delete destination *Destination*?

Return codes

- DEST_NAME_DOES_NOT_EXIST
Destination name does not exist
- CANNOT_CHANGE_EVENT_CONF_WITH_ALERTING_EVENTS
Cannot change event configuration while there are alerting events
Troubleshooting: Clear all alerting events before changing event configuration
- DEST_IS_PART_OF_DESTGROUP
Destination is part of a destination group and hence cannot be deleted
- DEST_APPEARS_IN_RULE
Destination 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 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

Field ID	Field output	Default position
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@yourcompany.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 event configuration while there are alerting events
Troubleshooting: Clear all alerting events before changing event configuration
- DEST_NAME_DOES_NOT_EXIST
Destination name does not exist

- DEST_NAME_IS_DESTGROUP_NAME
Destination name already exists as a destination group name
- DEST_NAME_ALREADY_EXISTS
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 ]
[ msgw=MSGGatewayName ] [ 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.
msgw	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
Destination name does not exist
- DEST_TEST_NOT_PERFORMED_SYSTEM_BUSY
Test of destination '*Destination Name*' not performed because the system is busy

Troubleshooting: Please wait a few seconds and try again

- GATEWAY_NAME_DOES_NOT_EXIST
Gateway name does not exist
- MSGWS_MUST_BE_SPECIFIED_FOR_DEST_TYPE
Destination must have SMS gateways
- MSGWS_NOT_ALLOWED_FOR_DEST_TYPE
Destination cannot have SMS gateways
- SMTPGWS_MUST_BE_SPECIFIED_FOR_DEST_TYPE
Destination must have SMTP gateways
- SMTPGWS_NOT_ALLOWED_FOR_DEST_TYPE
Destination cannot have SMTP gateways
- DEST_TEST_FAILED
Test of destination '*Destination Name*' failed
- SYSTEM_HAS_NO_SUCH_EXTERNAL_IP
The system has no such external IP address
- MODULE_CANNOT_SEND_MESSAGES
Selected module cannot send messages

Troubleshooting: Contact support

- ONLY_TECHNICIAN_CAN_REFER_TO_INTERNAL_EVENT_OBJECTS
Only technician can 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 ]
[ msggws=<MSGGW1[,MSGGW2]...|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.
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.
msggws	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@yourcompany.com
```

Output:

Command executed 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
Destination name does not exist
- AREA_CODE_MUST_BE_SPECIFIED_FOR_DEST_TYPE
Destination must have an area code
- AREA_CODE_NOT_ALLOWED_FOR_DEST_TYPE
Destination cannot have an area code
- CANNOT_CHANGE_EVENT_CONF_WITH_ALERTING_EVENTS
Cannot change event configuration while there are alerting events

Troubleshooting: Clear all alerting events before changing event configuration
- EMAIL_MUST_BE_SPECIFIED_FOR_DEST_TYPE
Destination must have an email address
- EMAIL_NOT_ALLOWED_FOR_DEST_TYPE
Destination cannot have an email address
- GATEWAY_NAME_APPEARS_TWICE
Gateway name appears twice in the list
- GATEWAY_NAME_DOES_NOT_EXIST
Gateway name does not exist
- NUMBER_MUST_BE_SPECIFIED_FOR_DEST_TYPE
Destination must have a number
- NUMBER_NOT_ALLOWED_FOR_DEST_TYPE
Destination cannot have a number
- SMSGWS_NOT_ALLOWED_FOR_DEST_TYPE
Destination cannot have SMS gateways
- SNMP_MANAGER_NOT_ALLOWED_FOR_DEST_TYPE
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
Destination cannot simultaneously have an email address and refer to a user
- DEST_CANNOT_HAVE_A_USER_AND_A_PHONE_NUMBER
Destination cannot simultaneously have a phone number address and refer to a user
- USER_PHONE_NUMBER_IS_NOT_DEFINED
User's phone number is not defined
- USER_NAME_DOES_NOT_EXIST
User name does not exist
- INTERNAL_DESTS_CANNOT_REFER_TO_USERS
Internal destinations cannot refer to users
- DEST_HEARTBEAT_DAYS_BUT_NO_HOUR
Destination heartbeat days specified with no heartbeat hour
- HTTPS_ADDRESS_NOT_ALLOWED_FOR_DEST_TYPE
Destination cannot have an https address
- PROXY_ADDRESS_NOT_ALLOWED_FOR_DEST_TYPE
Destination cannot have a proxy address
- SNMP_DESTS_CANNOT_REFER_TO_USERS
SNMP destinations cannot refer to users
- HTTPS_DESTS_CANNOT_REFER_TO_USERS
HTTPS destinations cannot refer to users
- USER_EMAIL_ADDRESS_IS_NOT_DEFINED
User's email address is not defined
- SMTPGWS_NOT_ALLOWED_FOR_DEST_TYPE
Destination cannot have SMTP gateways
- DAY_APPEARS_TWICE
The day 'Day' appears twice on the list.

Troubleshooting: Each day must appear at most once.
- SNMP_MANAGER_MUST_BE_SPECIFIED_FOR_DEST_TYPE
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
Destination must have an https address
 - DEST_TYPE_NOT_SUPPORTED
This type of destination is not supported
- Troubleshooting:** Contact support
- DOMAIN_DOESNT_EXIST
Domain does not exist.
 - USER_IS_NOT_IN_DESTINATION_DOMAINS
User must be included in the destination domains.
 - DESTINATION_IS_NOT_IN_RULE_DOMAINS
Destination must be included in the rule domains.
 - DESTINATION_IS_NOT_IN_DESTGROUP_DOMAINS
Destinations 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 executed 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
Destination group name does not exist
 - CANNOT_CHANGE_EVENT_CONF_WITH_ALERTING_EVENTS
Cannot change event configuration while there are alerting events
- Troubleshooting:** Clear all alerting events before changing event configuration
- DEST_NAME_DOES_NOT_EXIST
Destination name does not exist
 - DESTGROUP_MAX_DESTS_REACHED
Maximum number of destinations already defined in destination groups
 - DESTGROUP_ALREADY_INCLUDES_DEST
Destination group already includes destination name

- DESTINATION_IS_NOT_IN_DESTGROUP_DOMAINS
Destinations 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 executed 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
Maximum number of destination groups already defined
- DESTGROUP_NAME_ALREADY_EXISTS
Destination group name already exists
- DESTGROUP_NAME_IS_DEST_NAME
Destination group name already exists as a destination name
- DOMAIN_DOESNT_EXIST
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 executed 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
Destinations must be included in the destination group domains.
- DOMAIN_DOESNT_EXIST
Domain does not exist.
- DESTGROUP_IS_NOT_IN_RULE_DOMAINS
Destination groups must be included in the rule domains.
- DESTGROUP_NAME_DOES_NOT_EXIST
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

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 event configuration while there are alerting events
Troubleshooting: Clear all alerting events before changing event configuration
- DESTGROUP_APPEARS_IN_RULE
Destination Group appears in a Rule
Troubleshooting: To delete the destination group, first delete the rule.
- DESTGROUP_NAME_DOES_NOT_EXIST
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
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 executed 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
Destination group name does not exist
 - DEST_NAME_DOES_NOT_EXIST
Destination name does not exist
 - CANNOT_CHANGE_EVENT_CONF_WITH_ALERTING_EVENTS
Cannot change event configuration while there are alerting events
- Troubleshooting:** Clear all alerting events before changing event configuration
- DESTGROUP_DOES_NOT_INCLUDE_DEST
Destination group does not include 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 executed 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
Destination group name does not exist
 - DESTGROUP_NAME_ALREADY_EXISTS
Destination group name already exists
 - CANNOT_CHANGE_EVENT_CONF_WITH_ALERTING_EVENTS
Cannot change event configuration while there are alerting events
- Troubleshooting:** Clear all alerting events before changing event configuration
- DESTGROUP_NAME_IS_DEST_NAME
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

Name	Type	Description	Mandatory	Default
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 executed 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 technician can 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|smgw|smtpgw|target|volume|
cluster|ip_interface|ldap_conf|meta_data_object|
sync_schedule|user|user_group|ldap_server|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	Maximum number of events to list.	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

Name	Type	Description	Mandatory	Default
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.

Events are listed and sorted by time of creation, where the latest events are listed last. Events are listed by default in the user-readable textual form. Alternatively, the CLI option for comma-separated values can be used to generate output that can serve as input for other applications.

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.

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

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 15:16:22 Critical      WOULD_BE_EMERGENCY_SHUTDOWN
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
xiv_development  System has entered ON state.
                Storage Pool of size 171GB was created with name
                'pl_m'.
                An emergency shutdown has been detected, but UPS control
                is disabled.
xiv_development  Volume was created with name 'master' and size 17GB in
                Storage Pool with name 'pl_m'.
```

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 event code
Troubleshooting: Consult the manual for the list of event codes
- CANNOT_READ_EVENTS
Cannot read events.
Troubleshooting: Contact support
- DOMAIN_DOESNT_EXIST

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

User Category	Permission
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- EVENT_DOES_NOT_HAVE_THRESHOLDS
Event does not have thresholds

- EVENT_THRESHOLD_IS_ILLEGAL
Illegal value for event threshold

Troubleshooting: Event threshold values must be monotonic

- UNRECOGNIZED_EVENT_CODE
'String' is not a recognized event code

Troubleshooting: Consult the manual for the list of event codes

- LAST_EVENT_THRESHOLD_CANNOT_BE_DELETED
Events 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.

User Category	Permission	Condition
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

Generating a custom event

Use the **mm_event** command to generate a maintenance module event.

```
mm_event description=Description
[ severity=<INFORMATIONAL|WARNING|MINOR|MAJOR|CRITICAL> ]
category=Category
mm_data=AdditionalData
```

Parameters

Name	Type	Description	Mandatory	Default
description	String	Description of the event.	Y	N/A
severity	N/A	Severity of the event.	N	Informational
category	String	Category of the event.	Y	N/A
mm_data	String	Additional data for the event.	Y	N/A

This command can be used for either generating an event from a user application or host side software, or in order to test the event notifications procedures.

Example:

```
mm_event description="Description" category=Disk mm_data="Additional Data"
```

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 a maintenance module's heartbeat

Use the **mm_heartbeat** command to retrieve the heartbeat of a maintenance module.

```
mm_heartbeat data=HeartbeatData mm=ComponentId serial=SN part_number=PartNumber
total_memory=2048 free_memory=100 temperature=20 version=4.6 free_disk_tmp=200
free_disk_root=1000 free_disk_opt=900 free_disk_var=900
```

Parameters

Name	Type	Description	Mandatory
data	String	Heartbeat data.	Y
mm	N/A	Component ID of the maintenance module.	Y
serial	String	Serial number of the maintenance module.	Y
part_number	String	Part number of the maintenance module.	Y
version	String	Version of SW running on the maintenance module.	Y
total_memory	Integer	Total RAM of the maintenance module.	Y
free_memory	Integer	Free RAM on the maintenance module.	Y
temperature	Integer	Temperature of the maintenance module.	Y
free_disk_root	Integer	Free disk space on root partition of the maintenance module.	Y
free_disk_var	Integer	Free disk space on var partition of the maintenance module.	Y
free_disk_opt	Integer	Free disk space on opt partition of the maintenance module.	Y
free_disk_tmp	Integer	Free disk space on tmp partition of the maintenance module.	Y

Example:

```
mm_heartbeat data="" mm=1:MaintenanceModule:16 serial=serial part_number=pn
total_memory=2048 free_memory=100 temperature=20 version=4.6 free_disk_tmp=200
free_disk_root=1000 free_disk_opt=900 free_disk_var=900
```

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

- COMPONENT_DOES_NOT_EXIST
Component does not exist
- HEARTBEAT_DATA_TOO_LONG
Heartbeat data sent by the Maintenance Module is too long

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
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
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

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 event_clear command (see [Clearing alerting events](#)).

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 parameter.
- 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 executed 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
Maximum number of event rules already defined

- **EVENT_RULE_CANNOT_ESCALATE_TO_NON_ALERTING_RULES**
Event rule cannot escalate to non-alerting rule
- **Troubleshooting:** Alerting rule can only escalate to another escalating rule
- **DEST_APPEARS_TWICE**
Destination or destination group appears twice
- **EVENT_RULE_NAME_ALREADY_EXISTS**
Event rule name already exists
- **EVENT_RULE_NAME_DOES_NOT_EXIST**
Event rule name does not exist
- **NAME_IS_NEITHER_DEST_NOR_GROUP**
Name is neither the name of a destination group nor the name of a destination
- **ESCALATION_TIME_MUST_BE_LARGER_THAN_SNOOZE_TIME**
Escalation time must be larger than snooze time
- **RULE_MAX_DESTS_REACHED**
Maximum number of destinations and destination groups in a rule already defined
- **EVENT_RULE_MUST_HAVE_FILTER**
An alerting event rule must have a filter, either 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**
Escalation-only event rules must be alerting rules
- **TOO_MANY_EVENT_CODES**
A maximum of *Maximum* event codes can be specified
- **EVENT_CODE_APPEARS_TWICE**
Event code '*Code*' appears twice in the list

Troubleshooting: Each event code must appear at most once.

- **UNRECOGNIZED_EVENT_CODE**
'*String*' is not a recognized event code

Troubleshooting: Consult the manual for the list of event codes

- **EVENT_RULE_CANNOT_HAVE_A_CATEGORY**
A user event rule cannot have a category definition
- **DOMAIN_DOESNT_EXIST**
Domain does not exist.
- **DESTINATION_IS_NOT_IN_RULE_DOMAINS**
Destination must by included in the rule domains.
- **DESTGROUP_IS_NOT_IN_RULE_DOMAINS**
Destination groups must by included in the rule domains.
- **ESCALATION_RULE_NOT_IN_RULE_DOMAINS**
Escalation rule must belong to rule domains.
- **EVENT_RULE_MUST_NOT_HAVE_SNMP_DEST**
On current system configuration state (snmp_type = NONE) event rule must not have 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`
Event rule name does not exist
- `ESCALATION_ONLY_RULES_ARE_ALWAYS_ACTIVE`
Escalation-only event rules 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 event configuration while there are alerting events
Troubleshooting: Clear all alerting events before changing event configuration
- `EVENT_RULE_NAME_DOES_NOT_EXIST`
Event rule name does not exist
- `EVENT_RULE_USED_FOR_ESCALATION_CAN_NOT_BE_DELETED`
Event rule is an escalation rule of another event rule and thus 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

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

Example:

```
rule_list
```

Output:

Name	Minimum Severity	Event Code	Destinations
emergency_alerts	critical	all	john-cell,itstaff

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
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 event configuration while there are alerting events
Troubleshooting: Clear all alerting events before changing event configuration
- EVENT_RULE_NAME_ALREADY_EXISTS
Event rule name already exists
- EVENT_RULE_NAME_DOES_NOT_EXIST
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
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 executed 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**
Event rule name does not exist
- **CANNOT_CHANGE_EVENT_CONF_WITH_ALERTING_EVENTS**
Cannot change event configuration while there are alerting events
Troubleshooting: Clear all alerting events before changing 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**
Event rule cannot escalate to non-alerting rule
Troubleshooting: Alerting rule can only escalate to another escalating rule
- **DEST_APPEARS_TWICE**
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 also.
- **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**
Name is neither the name of a destination group nor the name of a destination
- **ESCALATION_TIME_MUST_BE_LARGER_THAN_SNOOZE_TIME**
Escalation time must be larger than snooze time
- **RULE_MAX_DESTS_REACHED**
Maximum number of destinations and destination groups in a rule already defined
- **EVENT_RULE_MUST_HAVE_FILTER**
An alerting event rule must have a filter, either event code or severity
- **CYCLIC_ESCALATION_RULES_DEFINITION**
Event rule escalation cannot be cyclic
- **EVENT_RULE_USED_FOR_ESCALATION_MUST_BE_ALERTING**
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**
Escalation-only event rules must be alerting rules
- **TOO_MANY_EVENT_CODES**
A maximum of *Maximum* event codes can be specified
- **EVENT_CODE_APPEARS_TWICE**
Event code 'Code' appears twice in the list

Troubleshooting: Each event code must appear at most once.

- UNRECOGNIZED_EVENT_CODE
'String' is not a recognized event code

Troubleshooting: Consult the manual for the list of event codes

- DOMAIN_DOESNT_EXIST
Domain does not exist.
- DESTINATION_IS_NOT_IN_RULE_DOMAINS
Destination must be included in the rule domains.
- DESTGROUP_IS_NOT_IN_RULE_DOMAINS
Destination groups must be included in the rule domains.
- ESCALATION_RULE_NOT_IN_RULE_DOMAINS
Escalation rule must belong to rule domains.
- EVENT_RULE_MUST_NOT_HAVE_SNMP_DEST
On current system configuration state (snmp_type = NONE) event rule must not have 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=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 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 define 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 of an area code and the local number. Both are specified when the destination is defined as described in the [Defining a new event notification destination](#) command.

The message text and destination numbers can be embedded into various parts of the email message: destination address, subject line, or email body. This command defines how email messages are formatted, and how the information of the specific SMS is arranged.

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. The information must be embedded into the following escape sequences:

- {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 smtpgw_prioritize command (see [Prioritizing SMTP gateways](#)). If needed, the user may define a specific SMTP gateway or gateways for sending email to this email-to-SMS gateway.

The system will try each SMS gateway, in the order specified in the smtpgw_prioritize command, until it successfully connects to one of them. The specific SMS destination can be associated with the specific SMS gateway (see [Defining a new event notification destination](#)).

Example:

```
smsgw_define smsgw=MSGW1
email_address={areacode}{number}@sms2emailserver.yourcompany.com
subject_line=SMS email_body={message}
```

Output:

Command executed 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
Maximum number of gateways already defined
- SMSGW_CANNOT_BE_DEFINED_WITHOUT_SMTPGW
SMS gateways cannot be defined if no SMTP gateways are defined
- GATEWAY_NAME_DOES_NOT_EXIST
Gateway name does not exist
- GATEWAY_NAME_APPEARS_TWICE
Gateway name appears twice in the list
- GATEWAY_NAME_ALREADY_EXISTS
Gateway name already exists

Deleting an SMS gateway

Use the **msgw_delete** command to delete an SMS gateway.

```
msgw_delete msgw=SMSSGatewayName
```

Parameters

Name	Type	Description	Mandatory
msgw	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:

```
msgw_delete msgw=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 event configuration while there are alerting events

Troubleshooting: Clear all alerting events before changing event configuration

- GATEWAY_NAME_DOES_NOT_EXIST
Gateway name does not exist
- GATEWAY_USED_BY_DESTINATION
Gateway is used by a destination

Listing SMS gateways

Use the **msgw_list** command to list SMS gateways.

```
msgw_list [ msgw=SMSSGatewayName ]
```

Parameters

Name	Type	Description	Mandatory	Default
msgw	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:

```
msgw_list
```

Output:

Name	Email Address	SMTP Gateways
MSGW1	{areacode} {number}@sms2emailserver.yourcompany.com	all
MSGW2	{areacode} {number}@sms2emailservice.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 **msgw_prioritize** command to set the priorities of the SMS gateways for sending SMS messages.

```
msgw_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:

```
msgw_prioritize order=MSGW1,MSGW2
```

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 event configuration while there are alerting events

Troubleshooting: Clear all alerting events before changing event configuration
- GATEWAY_NAME_APPEARS_TWICE
Gateway name appears twice in the list
- GATEWAY_NAME_DOES_NOT_EXIST
Gateway name does not exist
- GATEWAY_NAME_MISSING_FROM_LIST
Gateway name is missing from the list

Renaming an SMS gateway

Use the **msgw_rename** command to rename an SMS gateway.

```
msgw_rename msgw=SMSSGatewayName new_name=Name
```

Parameters

Name	Type	Description	Mandatory
msgw	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:

```
msgw_rename msgw=SMMSGW2 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 event configuration while there are alerting events

Troubleshooting: Clear all alerting events before changing event configuration
- GATEWAY_NAME_ALREADY_EXISTS
Gateway name already exists
- GATEWAY_NAME_DOES_NOT_EXIST
Gateway name does not exist

Updating an SMS gateway

Use the **msgw_update** command to update an SMS gateway.

```
msgw_update msgw=SMSSGatewayName [ email_address=email ]
[ subject_line=SubjectLineScheme ] [ email_body=EmailBodyScheme ]
[ smtpgw=<SMTPGW1[,SMTPGW2]...|ALL> ]
```

Parameters

Name	Type	Description	Mandatory	Default
msgw	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:

```
msgw_update msgw=SMSSGW1
email_address={areacode}{number}@sms2emailserver.yourcompany.com
subject_line=NextraSMS
email_body={message}
```

Output:

Command executed 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 event configuration while there are alerting events
Troubleshooting: Clear all alerting events before changing event configuration
- GATEWAY_NAME_APPEARS_TWICE
Gateway name appears twice in the list
- GATEWAY_NAME_DOES_NOT_EXIST
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_addresses	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.yourcompany.com
from_address=extra@yourcompany.com
reply_to_address=extraerrors@yourcompany.com
```

Output:

```
Command executed 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
Maximum number of gateways already defined
- GATEWAY_NAME_ALREADY_EXISTS
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 event configuration while there are alerting events
Troubleshooting: Clear all alerting events before changing event configuration
- GATEWAY_NAME_DOES_NOT_EXIST
Gateway name does not exist
- GATEWAY_USED_BY_DESTINATION
Gateway is used by a destination
- GATEWAY_USED_BY_SMS_GATEWAY
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.yourcompany.com	25	1
mailserver2	smtp.yourcompany.com	25	2

Access control

User Category	Permission	Condition
Storage administrator	Conditionally Allowed	Allowed, unless the internal parameter is specified.

User Category	Permission	Condition
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 prioritize 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, XIV 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 event configuration while there are alerting events
Troubleshooting: Clear all alerting events before changing event configuration
- GATEWAY_NAME_APPEARS_TWICE
Gateway name appears twice in the list
- GATEWAY_NAME_DOES_NOT_EXIST
Gateway name does not exist
- GATEWAY_NAME_MISSING_FROM_LIST
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 event configuration while there are alerting events
Troubleshooting: Clear all alerting events before changing event configuration
- GATEWAY_NAME_ALREADY_EXISTS
Gateway name already exists
- GATEWAY_NAME_DOES_NOT_EXIST
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_addresses	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.yourcompany.com  
from_address=nextra@yourcompany.com  
reply_to_address=nextraerrors@yourcompany.com
```

Output:

```
Command executed 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

User Category	Permission	Condition
Read-only users	Disallowed	N/A
Technicians	Allowed	N/A

Return codes

- GATEWAY_NAME_DOES_NOT_EXIST
Gateway name does not exist
- CANNOT_CHANGE_EVENT_CONF_WITH_ALERTING_EVENTS
Cannot change event configuration while there are alerting events
Troubleshooting: Clear all alerting events before changing event configuration
- FROM_ADDRESS_NOT_DEFINED
Neither the gateway's From Address nor the default From Address is defined

Generating an XMPNS user control event

Use `xmpns_user_config_set` command to generate an **XMPNS_USER_CONTROL** event.

```
xmpns_user_config_set action=Action
```

Parameters

Name	Type	Description	Mandatory
action	String	Action code text.	Y

This command generates an XMPNS_USER_CONTROL event which includes the `action_code` text in the event's description field. The current logged-in username is also added to the action string sent in the description field after validating that the user exists in the given system.

Example:

```
xmpns_user_config_set action="registration;device_token=aa23d1234;pns=gcm"
```

Output:

```
Command executed 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

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_code` 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 executed 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.

- [Adding Ethernet ports to IP interfaces](#)
Use the `ipinterface_add_port` command to add an Ethernet port to the link aggregation group of an IP interface.
- [Creating a new IP interface](#)
Use the `ipinterface_create` command to create a new IP interface for iSCSI.
- [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.
- [Showing the status and configuration of Ethernet ports](#)
Use the `ipinterface_list_ports` command to list all Ethernet ports together with their configuration and status.
- [Removing Ethernet ports from IP interfaces](#)
Use the `ipinterface_remove_port` command to remove an Ethernet port from the link aggregation group of an IP interface.
- [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 IP](#)
Use the `ipinterface_run_traceroute6` command to test connectivity to a remote IP node using the ICMP trace-route mechanism.
- [Updating an IP interface](#)
Use the `ipinterface_update` command to update the configuration of an IP interface.
- [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 all or specific 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 the status of all support centers](#)
Use the `support_center_status` command to list information about all defined support centers.
- [Defining the automatic connection to a support center](#)
Use the `support_center_config` command to configure the automatic connection to a support center.
- [Listing the configuration of automatically connected support centers](#)
Use the `support_center_config_list` command to display the configuration of automatically connected support centers.
- [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 groups](#)
Use the `ip_access_group_address_list` command to list IP access group addresses.

Adding Ethernet ports to IP interfaces

Use the `ipinterface_add_port` command to add an Ethernet port to the link aggregation group of an IP interface.

```
ipinterface_add_port ipinterface=IPInterfaceName port=P
```

Parameters

Name	Type	Description	Mandatory
ipinterface	Object name	IP interface to which the port is to be added.	Y
port	Integer	Number of the port to be added to the group.	Y

The specified port is added to the link aggregation group of the specified IP interface.

Ports defined as a link aggregation group must be connected to the same Ethernet switch, and a parallel link aggregation group must be defined on that Ethernet switch.

The module is not provided, as it must be the module of the other ports of the interface.

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`
IP Interface name does not exist.
- `ILLEGAL_PORT_NUMBER`
Port number is out of range.
- `PORT_ALREADY_IN_INTERFACE`
Port is already part of the specified IP interface.
- `PORT_IS_USED_IN_ANOTHER_IP_INTERFACE`
One of the physical ports specified is already assigned to an IP interface.
- `COMMAND_NOT_ALLOWED_ON_MANAGEMENT_OR_VPN_INTERFACE`
Operation is not allowed on management or VPN IP interface.
- `IPINTERFACE_NOT_AGGREGATED`
Ports cannot be added to a non-aggregated IP interface.

Creating a new IP interface

Use the `ipinterface_create` command to create a new IP interface for iSCSI.

```
ipinterface_create ipinterface=IPInterfaceName address=Address netmask=NetworkMask
[ gateway=DefaultGateway ] [ mtu=MTU ] module=ModuleNumber port=PortNumber
[ speed=<auto|10mb|100mb|1000mb|1gb|2500mb|2.5gb|10000mb|10gb> ]
```

Parameters

Name	Type	Description	Mandatory	Default

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. This parameter is optional.	N	None
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.
module	N/A	Component identifier (rack and module) of the module containing Ethernet ports.	Y	N/A
port	Integer	Port Number	Y	N/A
speed	Enumeration	Interface's speed, either automatic or explicit. An explicit speed turns off auto-negotiation.	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**
IP Interface name already used
- **ILLEGAL_PORT_NUMBER**
Port number is out of range
- **PORT_IS_USED_IN_ANOTHER_IP_INTERFACE**
One of the physical ports specified is already assigned to an IP Interface
- **PORT_REPEATS_TWICE**
Port list contains the same value more than once
- **FORCE_NO_AGGREGATION_ALLOWED_FOR_SINGLE_PORT_ONLY**
More than one port specified for non-aggregated IP Interface
- **IP_ADDRESS_ALREADY_USED_IN_ANOTHER_INTERFACE**
IP address is already assigned to another interface
- **IPADDRESS_AND_GATEWAY_ARE_NOT_ON_SAME_SUBNET**
IP address specified for the default gateway is not in the subnet of the IP Interface
- **MTU_TOO_LARGE**
Specified MTU value is too large
- **BAD_PORTS_FORMAT**
Port list should be a comma separated list of positive integers
- **ILLEGAL_COMPONENT_ID**
Component ID is illegal
- **TOO_MANY_PORTS_IN_AGGREGATION_GROUP**
Too many physical ports for one IP interface
- **ILLEGAL_IPADDRESS**
Illegal IP address was entered
- **DUPLICATE_IPADDRESSES**
Duplicate IP addresses were specified
- **ILLEGAL_GATEWAY_IPADDRESS**
Illegal IP address was specified for default gateway

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
IP Interface name does not exist
- COMMAND_NOT_ALLOWED_ON_MANAGEMENT_OR_VPN_INTERFACE
Operation is not allowed on Management or VPN IP Interface
- IPINTERFACE_HAS_CONNECTIVITY
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  9.151.154.239  255.255.248.0  9.151.159.254
-----
```

Cont.:

```
IPv6 Gateway  MTU      Module      Port  IP access group name
-----
              1500    1:Module:12
```

Field ID	Field output	Default position
name	Name	1
type	Type	2
address	IP Address	3
netmask	Network Mask	4
gateway	Default Gateway	5
address6	IPv6 Address	6
gateway6	IPv6 Gateway	7
mtu	MTU	8
module	Module	9
port	Port	10
speed	Speed	N/A
access_group	IP access group name	11

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=ModuleNumber ]
```

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:bf8:2000:5159:42f2:e9ff:feaf:ccb2/64
management  Management      9.151.154.239/21
management  Management      fe80::42f2:e9ff:feaf:ccb2/64
```

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

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Showing the status and configuration of Ethernet ports

Use the **ipinterface_list_ports** command to list all Ethernet ports together with their configuration and status.

ipinterface_list_ports

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

Removing Ethernet ports from IP interfaces

Use the **ipinterface_remove_port** command to remove an Ethernet port from the link aggregation group of an IP interface.

```
ipinterface_remove_port ipinterface=IPInterfaceName port=P
```

Parameters

Name	Type	Description	Mandatory
ipinterface	Object name	IP interface from which the port is to be removed.	Y
port	Integer	Number of the port to be removed from the group.	Y

This command removes the specified port from the link aggregation group of the specified IP interface. The module does not need to be specified, because it is the same module as the other ports of the IP interface.

The last port of the IP interface cannot be removed.

If the IP interface must be moved to a different module, first delete the interface and then recreate it. This command cannot be applied to management or VPN interfaces.

Access control

User Category	Permission
Storage administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Operations administrator	Allowed
Host side accelerator client	Disallowed

Return codes

- **IPINTERFACE_DOES_NOT_EXIST**
IP interface name does not exist.

- **ILLEGAL_PORT_NUMBER**
Port number is out of range.
- **PORT_NOT_IN_INTERFACE**
Port is not part of the specified IP interface.
- **COMMAND_NOT_ALLOWED_ON_MANAGEMENT_OR_VPN_INTERFACE**
Operation is not allowed on management or VPN IP interface.
- **IPINTERFACE_LAST_PORT**
Last port in IP interface cannot be removed.

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**
IP Interface name does not exist
- **IPINTERFACE_EXISTS**
IP Interface name already used
- **COMMAND_NOT_ALLOWED_ON_MANAGEMENT_OR_VPN_INTERFACE**
Operation is not allowed on 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 | localipaddress6=IPv6address
```

Parameters

Name	Description	Mandatory
localipaddress	IP address of the IP interface for which the ARP database should be printed.	N
localipaddress6	IPv6 address of the IP interface for which the ARP database should be printed.	N

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

User Category	Permission
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 matches given criteria
- MORE_THAN_ONE_IP_INTERFACE_MATCHES
More than one IP Interface matches given 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
```

Parameters

Name	Description	Mandatory
localipaddress	IP address of the IP interface for which the traceroute command is run.	Y
remote	IP 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 matches given criteria
- MORE_THAN_ONE_IP_INTERFACE_MATCHES
More than one IP Interface matches given criteria

Testing the traceroute to a remote IP

Use the **ipinterface_run_traceroute6** command to test connectivity to a remote IP node 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 matches given criteria
- MORE_THAN_ONE_IP_INTERFACE_MATCHES
More than one IP Interface matches given 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 ]  
[ mtu=MTU ] [ access_group=IPAccessGroupName ]
```

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.	N	Keep unchanged.
access_group	Object name	The name of the IP access group used for IP filtering.	N	Keep unchanged.

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.

Example:

```
ipinterface_update ipinterface=management
```

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

- **IPINTERFACE_DOES_NOT_EXIST**
IP Interface name does not exist
- **IP_ADDRESS_ALREADY_USED_IN_ANOTHER_INTERFACE**
IP address is already assigned to another interface
- **IPADDRESS_AND_GATEWAY_ARE_NOT_ON_SAME_SUBNET**
IP address specified for the default gateway is not in the subnet of the IP Interface
- **IPINTERFACE_MANAGEMENT_DIFFERENT_SUBNET**
IP addresses management modules must all be in the same subnet
- **IPINTERFACE_MANAGEMENT_MISSING_IPS**
Number of IP addresses specified is less than the number of management modules
- **IPINTERFACE_MANAGEMENT_TOO_MANY_IPS**
Number of IP addresses specified is larger than the number of management modules
- **MTU_TOO_LARGE**
Specified MTU value is too large
- **ILLEGAL_IPADDRESS**
Illegal IP address was entered
- **DUPLICATE_IPADDRESSES**
Duplicate IP addresses were specified
- **ILLEGAL_GATEWAY_IPADDRESS**
Illegal IP address was specified for default gateway
- **ILLEGAL_IPV6ADDRESS**
Illegal IPv6 address was entered
- **DUPLICATE_IPV6ADDRESSES**
Duplicate IPv6 addresses were specified
- **ILLEGAL_GATEWAY_IPV6_ADDRESS**
Illegal IPv6 address was specified for default gateway
- **IPV6ADDRESS_AND_GATEWAY_ARE_NOT_ON_SAME_SUBNET**
IPv6 address specified for the default gateway is not in the subnet of the IP Interface
- **IPV6_ADDRESS_ALREADY_USED_IN_ANOTHER_INTERFACE**
IPv6 address is already assigned to another interface
- **IPINTERFACE_MANAGEMENT_MISSING_IPV6S**
Number of IPv6 addresses specified is less than the number of management modules
- **IPINTERFACE_MANAGEMENT_TOO_MANY_IPV6S**
Number of IPv6 addresses specified is larger than the number of management modules
- **IPINTERFACE_MANAGEMENT_DIFFERENT_IPV6_SUBNET**
IPv6 addresses management modules must all be in the same subnet
- **IP_ACCESS_GROUP_DOES_NOT_EXIST**
IP access group with such name doesn't exist
- **IP_ACCESS_INVALID_INTERFACE_TYPE**
IP filtering is applied to an invalid interface (should be management or VPN)

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

Name	Type	Description	Mandatory	Default
passkey	N/A	Secret password.	N	N/A
certificate	N/A	The content of a <code>.pem</code> file, with asterisks (*) instead of newlines. In Windows, drag-and-drop the <code>.pem</code> file from the Windows Explorer to the appropriate location in the XCLI session window; the content will be added automatically.	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

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 IPsec connections is already configured
- **IPSEC_UNSUPPORTED_FOR_ISCSI**
IPsec is unsupported for iSCSI ports
- **SSL_CERTIFICATE_CHAIN_EMPTY**
No certificates found in input.
- **SSL_CERTIFICATE_HAS_EXPIRED**
SSL certificate has expired.
- **SSL_CERTIFICATE_INVALID_FORMAT**
SSL certificate format is invalid or corrupted.
- **SSL_CERTIFICATE_ISSUER_NOT_FOUND**
SSL certificate issuer not found in certificate chain.
- **SSL_CERTIFICATE_NOT_YET_VALID**
SSL certificate is not yet valid.
- **SSL_CERTIFICATE_VERIFICATION_FAILED**
SSL certificate chain verification failed.
- **SSL_CERTIFICATE_VERIFICATION_INTERNAL_ERROR**
SSL certificate verification has failed because of 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 content of a .pem file, with asterisks (*) instead of newlines. In Windows, drag-and-drop the .pem file from the Windows Explorer to the appropriate location in the XCLI session window; the content will be added automatically.	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.

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 unsupported for iSCSI ports
- SSL_CERTIFICATE_CHAIN_EMPTY
No certificates found in input.
- SSL_CERTIFICATE_HAS_EXPIRED
SSL certificate has expired.
- SSL_CERTIFICATE_INVALID_FORMAT
SSL certificate format is invalid or corrupted.
- SSL_CERTIFICATE_ISSUER_NOT_FOUND
SSL certificate issuer not found in certificate chain.
- SSL_CERTIFICATE_NOT_YET_VALID
SSL certificate is not yet valid.
- SSL_CERTIFICATE_VERIFICATION_FAILED
SSL certificate chain verification failed.
- SSL_CERTIFICATE_VERIFICATION_INTERNAL_ERROR
SSL certificate verification has failed because of 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:

```
xcli.py 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 all or specific 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

Name	Type	Description	Mandatory	Default
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, if specified, must be lower than the regular timeout

- **MODULE_HAS_NO_SUPPORT_CENTER_PORT**
The specified module does not have a port from which Support Center can connect
- **NO_MODULE_WITH_SUPPORT_CENTER_PORT**
No module has a port from which Support Center can connect
- **REMOTE_SUPPORT_CLIENT_ALREADY_RUNNING**
The Remote Support Client is already running

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:

```
xcli.py support_center_define support_center=somewhere address=1.1.1.1
```

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

- **MAX_SUPPORT_CENTERS_DEFINED**
Maximum number of support centers is already defined.

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

User Category	Permission
Technicians	Allowed

Warnings

- ARE_YOU_SURE_TO_DELETE_THE_SUPPORT_CENTER
Are you sure you want to delete Support Center?.

Return codes

- SUPPORT_CENTER_NOT_DEFINED
Support Center is not defined.
- CANNOT_DELETE_WHILE_SUPPORT_CENTER_IS_RUNNING
Support Center is running. Disconnect 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 the status of all support centers

Use the **support_center_status** command to list information about all 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
```

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

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 the automatic connection to a support center

Use the **support_center_config** command to configure the automatic connection to a support center.

```
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 the automatic connection to the support center.	N	Module with first management port

Name	Type	Description	Mandatory	Default
connect_through_h_module2	N/A	The second module from which to establish the automatic connection to the support center.	N	Module with second management port
connect_through_h_module3	N/A	The third module from which to establish the automatic connection to the support center.	N	Module with second 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. Format: string, must be 6-12 alphanumeric characters, and is case-insensitive.	N	none

Example:

```
xcli.py 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 Support Center can connect
- **NO_MODULE_WITH_SUPPORT_CENTER_PORT**
No module has a port from which Support Center can connect

Listing the configuration of automatically connected support centers

Use the **support_center_config_list** command to display the configuration of automatically connected support centers.

```
support_center_config_list
```

Example:

```
support_center_config_list
```

Output:

Enable Auto Conn	First Module	Second Module	Third Module
yes	1	2	-1

Access control

User Category	Permission
Storage administrator	Allowed
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
IP access group with such name already exists
- IP_ACCESS_MAXIMUM_NUMBER_OF_GROUPS_IS_REACHED
Reached maximum number of IP access groups

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=172.30.214.202
```

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
IP access group with such name doesn't exist
- IP_ACCESS_ADDRESS_IS_NOT_VALID
The given address is not valid

- IP_ACCESS_ADDRESS_IS_NOT_IN_GROUP
The given address isn't 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=172.30.214.202
```

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
IP access group with such name doesn't 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
Reached the maximum number of addresses in the IP access group

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
IP access group with such name doesn't 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

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
IP access group with such name doesn't exist
- IP_ACCESS_GROUP_ALREADY_EXISTS
IP access group with such 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

Field ID	Field output	Default position
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 groups

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	192.168.1.10
IPAccessGroup1	172.30.214.202

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
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 the 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 the certificate signing request
Troubleshooting: Generate a certificate signing request specifying a correct subject (e.g., '/C=US/CN=IBM')
- **CERTIFICATE_NAME_DOES_NOT_EXIST**
Certificate name was not found
Troubleshooting: Choose 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 the certificate signing request
Troubleshooting: Generate a certificate signing request specifying a correct subject (e.g., '/C=US/CN=IBM')
- **FAILED_CREATING_PRIVATE_KEY**
Failed creating private key
- **CERTIFICATE_NAME_ALREADY_EXIST**
Certificate with same name already exist
Troubleshooting: Choose a different name
- **CERTIFICATE_CONTAINER_FULL**
Can't add more certificates, the maximum already defined

Deleting the 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 certificate?

Return codes

- CERTIFICATE_NAME_DOES_NOT_EXIST
Certificate name was not found

Troubleshooting: Choose a different name
- DEFAULT_CERTIFICATE_CANNOT_BE_DELETED
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
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**
Certificate with same name already exist
Troubleshooting: Choose a different name
- **CERTIFICATE_NAME_DOES_NOT_EXIST**
Certificate name was not found
Troubleshooting: Choose 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 content of signed certificate in .pem file format. Asterisks (*) can be used instead of newlines. In Windows, drag-and-drop the .pem file from the Windows Explorer to the appropriate location in the XCLI session window; the content will be added automatically.	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](#).

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 other certificate.
Troubleshooting: Edit the certificate used by service before
- **PRIVATE_KEY_ALREADY_HAS_OTHER_CERTIFICATE**
The private key matching this certificate already has other certificate
Troubleshooting: Use the **pki_update** command if you want to replace the certificate
- **CERTIFICATE_KEY_WAS_NOT_FOUND**
Failed to set certificate
Troubleshooting: Check the parameters
- **SSL_CERTIFICATE_CHAIN_EMPTY**
No certificates found in input.
- **SSL_CERTIFICATE_HAS_EXPIRED**
SSL certificate has expired.
- **SSL_CERTIFICATE_INVALID_FORMAT**
SSL certificate format is invalid or corrupted.
- **SSL_CERTIFICATE_ISSUER_NOT_FOUND**

SSL certificate issuer not found in certificate chain.

- **SSL_CERTIFICATE_NOT_YET_VALID**
SSL certificate is not yet valid.
- **SSL_CERTIFICATE_VERIFICATION_FAILED**
SSL certificate chain verification failed.
- **SSL_CERTIFICATE_VERIFICATION_INTERNAL_ERROR**
SSL certificate verification has failed because of 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: <code>base64 -w0 myCert.pfx</code>	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](#).

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 other certificate.
Troubleshooting: Edit the certificate used by service before
- **CANNOT_VALIDATE_PKCS12_FILE**
Failed validating PKCS#12 file.
Troubleshooting: Check PKCS#12 file content is encoded to base64, and the password is OK.
- **DEFAULT_CERTIFICATE_ALREADY_EXIST**
Other default certificate already exist.
Troubleshooting: Remove the default certificate, or make it not default.
- **CERTIFICATE_NAME_ALREADY_EXIST**
Certificate with same name already exist
Troubleshooting: Choose a different name
- **BAD_BASE64_DATA**
Data cannot be decoded as base-64 data.
- **FAILED_GETTING_PRIVATE_KEY_FINGERPRINT**
Failed getting private key fingerprint.

- **FAILED_ENCRYPTING_PRIVATE_KEY**
Failed encrypting private key.
- **CERTIFICATE_CONTAINER_FULL**
Can't add more certificates, the maximum already defined

Troubleshooting: Delete certificate

- **SSL_CERTIFICATE_CHAIN_EMPTY**
No certificates found in input.
- **SSL_CERTIFICATE_HAS_EXPIRED**
SSL certificate has expired.
- **SSL_CERTIFICATE_INVALID_FORMAT**
SSL certificate format is invalid or corrupted.
- **SSL_CERTIFICATE_ISSUER_NOT_FOUND**
SSL certificate issuer not found in certificate chain.
- **SSL_CERTIFICATE_NOT_YET_VALID**
SSL certificate is not yet valid.
- **SSL_CERTIFICATE_VERIFICATION_FAILED**
SSL certificate chain verification failed.
- **SSL_CERTIFICATE_VERIFICATION_INTERNAL_ERROR**
SSL certificate verification has failed because of internal system error.

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 certificate.
- **KEY_HAS_NO_CERTIFICATE**
The key has no signed certificate defined.
- **CERTIFICATE_NAME_DOES_NOT_EXIST**
Certificate name was not found

Troubleshooting: Choose 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=SigendCertificate ]
```

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 other certificate.
Troubleshooting: Edit the certificate used by service before
- NO_PKI_UPDATE_PARAMETERS_SPECIFIED
No parameters were specified for update
- CERTIFICATE_DOES_NOT_MATCH_PRIVATE_KEY
Certificate does not match private key
Troubleshooting: Use other certificate.
- CANNOT_SET_SERVICES_BEFORE_SETTING_CERTIFICATE
Can't set services before certificate.
Troubleshooting: Set certificate first.
- DEFAULT_CERTIFICATE_ALREADY_EXIST
Other default certificate already exist.
Troubleshooting: Remove the default certificate, or make it not default.
- CERTIFICATE_KEY_WAS_NOT_FOUND
Failed to set certificate
Troubleshooting: Check the parameters
- CERTIFICATE_NAME_DOES_NOT_EXIST
Certificate name was not found
Troubleshooting: Choose a different name
- SSL_CERTIFICATE_CHAIN_EMPTY
No certificates found in input.
- SSL_CERTIFICATE_HAS_EXPIRED
SSL certificate has expired.
- SSL_CERTIFICATE_INVALID_FORMAT
SSL certificate format is invalid or corrupted.
- SSL_CERTIFICATE_ISSUER_NOT_FOUND
SSL certificate issuer not found in certificate chain.
- SSL_CERTIFICATE_NOT_YET_VALID
SSL certificate is not yet valid.

- **SSL_CERTIFICATE_VERIFICATION_FAILED**
SSL certificate chain verification failed.
- **SSL_CERTIFICATE_VERIFICATION_INTERNAL_ERROR**
SSL certificate verification has failed because of internal system error.

Compression commands

This section describes the command-line interface (CLI) for data compression.

- **Activating compression**
Use the **system_compression_enable** command to activate the compression feature and allocate the necessary system resources.
- **Deactivating compression**
Use the **system_compression_disable** command to deactivate compression on XIV modules and release the system resources needed for RtC (CPU cores and memory).
- **Displaying the state of the estimated compression savings feature**
Use the **system_comprestate_get** command to display the operational state (enabled or disabled) of the estimated compression savings feature.
- **Controlling the continuous display of the estimated compression savings**
Use the **system_comprestate_set** command to activate or deactivate the continuous display of the estimated compression savings on the volume.

Activating compression

Use the **system_compression_enable** command to activate the compression feature and allocate the necessary system resources.

system_compression_enable

Compression can be activated if all of the following conditions are true:

- system state is On
- redundancy status is Full
- there is no failed or phased out data service (that is, cache node) in the system
- no hot upgrade is in progress
- compression is disabled.

Example:

```
xcli -u -c XIV1 system_compression_enable
```

Output:

Command executed 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_ENABLE_COMPRESSION**
Using XIV compression requires an additional license. Are you sure you want to enable compression on this system?

Return codes

- **COMPRESSION_IS_ACTIVE**
Compression is already enabled.
- **COMMAND_IS_NOT_VALID_IN_CURRENT_SYSTEM_STATE**
The requested command cannot be invoked in the current system state
- **COMPRESSION_ENABLE_IN_PROGRESS**
Compression feature is already in enabling process
- **COMPRESSION_DISABLE_IN_PROGRESS**
Compression feature is already in disabling process
- **COMPRESSION_ENABLE_REQUIRES_ALL_DATA_SERVICES**

Compression feature enabling requires all data services to be valid and system to be fully redundant

- **NO_SPACE**
The system does not have enough free space for the requested Storage Pool size
- **OPERATION_NOT_ALLOWED_WHEN_SYSTEM_IS_NON_REDUNDANT**
Operation is not allowed on non redundant system.

Deactivating compression

Use the **system_compression_disable** command to deactivate compression on XIV modules and release the system resources needed for RtC (CPU cores and memory).

system_compression_disable

Compression can be deactivated if all of the following conditions are true:

- compression is currently enabled
- there are no compressed volumes in the system
- the queue of volumes being or waiting to be compressed is empty
- no hot upgrade is in progress.

Example:

```
xcli -u -c XIV1 system_compression_disable
```

Output:

Command executed 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

- **COMPRESSION_IS_NOT_ACTIVE**
Compression was not enabled.
- **COMPRESSED_VOLUME_EXISTS**
Compression can't be disabled since compressed volume(s) exist.
- **VOLUME_TRANSFORM_QUEUE_NOT_EMPTY**
Compression can't be disabled since volume transform queue is not empty.
- **COMMAND_IS_NOT_VALID_IN_CURRENT_SYSTEM_STATE**
The requested command cannot be invoked in the current system state
- **COMPRESSION_ENABLE_IN_PROGRESS**
Compression feature is already in enabling process
- **COMPRESSION_DISABLE_IN_PROGRESS**
Compression feature is already in disabling process
- **COMPRESSION_RECOVERY_IN_PROGRESS**
Operation cannot be completed because another recovery process is in progress

Displaying the state of the estimated compression savings feature

Use the **system_comprestimate_get** command to display the operational state (enabled or disabled) of the estimated compression savings feature.

system_comprestimate_get

Example:

```
system_comprestimate_get
```

Output:

Disabled

Access Control:

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Allowed
Read-only users	Allowed
Technicians	Disallowed

Controlling the continuous display of the estimated compression savings

Use the **system_compreestimate_set** command to activate or deactivate the continuous display of the estimated compression savings on the volume.

```
system_compreestimate_set mode=<enabled|disabled>
```

Parameters

Name	Type	Description	Mandatory
mode	Enumeration	Defines whether to activate or deactivate the continuous estimation of volume compression savings.	Y

This command is used to activate or deactivate the continuous display of the estimated compression savings of volumes in the system. The default mode is set to enabled.

Example:

```
system_compreestimate_set mode=enabled
```

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

- COMPRESTIMATION_IS_ENABLED
Compreestimate is enabled.
- COMPRESTIMATION_IS_DISABLED
Compreestimate is disabled.

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 the configured InfiniBand switches](#)
Use the **switch_list** command to list the configured InfiniBand switches.

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
saved_info.last_state	State	6
saved_info.port_down_reason	Failure Reason	7
currently_functioning	Currently Functioning	N/A
requires_service	Requires Service	N/A
service_reason	Service Reason	N/A
sm_port_in_test	SM Port in Test	N/A
sm_port_is_master	SM Port is Master	N/A
sm_port_should_have_state	SM Port Target State	N/A
sm_port_status	SM Port Status	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`

Field ID	Field output	Default position
component_id	Switch	1
status	Status	2
sw_mgmt_status	Management Status	3
num_of_down_ports	Down Ports	4
last_succ_monitoring	Last Monitoring Time	5
ps1	Power Supply #1	6
ps2	Power Supply #2	7
prob_fans_num	Problematic Fans	8
prob_temp_num	Problematic Temperatures	9
prob_volt_num	Problematic Voltages	10
currently_functioning	Currently Functioning	N/A
temp_is4	IS4 Temperature	N/A
temp_ps_ambient	PS Ambient Temperature	N/A
temp_is4_ambient	IS4 Ambient Temperature	N/A
temp_board	Board Temperature	N/A
fans_rpm.0	Fan #1 RPM	N/A
fans_rpm.1	Fan #2 RPM	N/A

Field ID	Field output	Default position
fans_rpm.2	Fan #3 RPM	N/A
fans_rpm.3	Fan #4 RPM	N/A
voltage.0	Expected Voltage #1	N/A
expected_voltage.0	Expected Voltage #1	N/A
voltage.1	Expected Voltage #2	N/A
expected_voltage.1	Expected Voltage #2	N/A
voltage.2	Expected Voltage #3	N/A
expected_voltage.2	Expected Voltage #3	N/A
voltage.3	Expected Voltage #4	N/A
expected_voltage.3	Expected Voltage #4	N/A
voltage.4	Expected Voltage #5	N/A
expected_voltage.4	Expected Voltage #5	N/A
voltage.5	Expected Voltage #6	N/A
expected_voltage.5	Expected Voltage #6	N/A
voltage.6	Expected Voltage #7	N/A
expected_voltage.6	Expected Voltage #7	N/A
voltage.7	Expected Voltage #8	N/A
expected_voltage.7	Expected Voltage #8	N/A
mgmt_guid	Management GUID	N/A
fabric_guid	Fabric GUID	N/A
curr_vpd.chassis_type	Current VPD - Chassis type	N/A
curr_vpd.mgmt_type	Current VPD - Management type	N/A
curr_vpd.cpu_type	Current VPD - CPU type	N/A
curr_vpd.chassis_pn	Current VPD - Chassis P/N	N/A
curr_vpd.mgmt_pn	Current VPD - Management P/N	N/A
curr_vpd.cpu_pn	Current VPD - CPU P/N	N/A
curr_vpd.chassis_sn	Current VPD - Chassis S/N	N/A
curr_vpd.mgmt_sn	Current VPD - Management S/N	N/A
curr_vpd.cpu_sn	Current VPD - CPU S/N	N/A
curr_vpd.asic_fw_version	Current ASIC firmware version	N/A
curr_vpd.mgmt_fw_version	Current Management firmware version	N/A
prev_vpd.chassis_type	Previous VPD - Chassis type	N/A
prev_vpd.mgmt_type	Previous VPD - Management type	N/A
prev_vpd.cpu_type	Previous VPD - CPU type	N/A
prev_vpd.chassis_pn	Previous VPD - Chassis P/N	N/A
prev_vpd.mgmt_pn	Previous VPD - Management P/N	N/A
prev_vpd.cpu_pn	Previous VPD - CPU P/N	N/A
prev_vpd.chassis_sn	Previous VPD - Chassis S/N	N/A
prev_vpd.mgmt_sn	Previous VPD - Management S/N	N/A
prev_vpd.cpu_sn	Previous VPD - CPU S/N	N/A
prev_vpd.asic_fw_version	Previous ASIC firmware version	N/A
prev_vpd.mgmt_fw_version	Previous Management firmware version	N/A
initial_vpd.chassis_type	Initial VPD - Chassis type	N/A
initial_vpd.mgmt_type	Initial VPD - Management type	N/A
initial_vpd.cpu_type	Initial VPD - CPU type	N/A
initial_vpd.chassis_pn	Initial VPD - Chassis P/N	N/A
initial_vpd.mgmt_pn	Initial VPD - Management P/N	N/A
initial_vpd.cpu_pn	Initial VPD - CPU P/N	N/A
initial_vpd.chassis_sn	Initial VPD - Chassis S/N	N/A
initial_vpd.mgmt_sn	Initial VPD - Management S/N	N/A
initial_vpd.cpu_sn	Initial VPD - CPU S/N	N/A
initial_vpd.asic_fw_version	Initial ASIC firmware version	N/A
initial_vpd.mgmt_fw_version	Initial Management firmware version	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

User Category	Permission
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](#)
Use the **domain_list** command to 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.
- [Moving a grouped pool from one domain to another](#)
Use the **domain_move_gp** 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 executed 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
User group name does not exist
- CLUSTER_BAD_NAME

Cluster name does not exist

- HOST_BAD_NAME
Host name does not exist
- HOST_BELONGS_TO_CLUSTER
Host is part of 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 executed 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
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
Cluster name does not exist
- HOST_BAD_NAME
Host name does not exist
- HOST_BELONGS_TO_CLUSTER
Host is part of 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

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
Host name does not exist
- CLUSTER_BAD_NAME
Cluster name does not exist
- USER_GROUP_NAME_DOES_NOT_EXIST
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.	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 content of a .pem file, with asterisks (*) instead of newlines. In Windows, drag-and-drop the .pem file from the Windows Explorer to the appropriate location in the XCLI session window; the content will be added automatically.	N	no certificate
port	Integer	The port number.	N	389
secure_port	Integer	The secure port number.	N	636

Example:

```
ldap_add_server fqdn=ldap.example.com address=1.2.3.4
```

Output:

```
Command executed 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_LDAP_SERVERS_REACHED**
Maximum number of LDAP servers already defined
- **ADDRESS_CURRENTLY_ASSOCIATED_WITH_ANOTHER_LDAP_SERVER**
The specified IP address is currently associated with another LDAP server
- **LDAP_SERVER_EXISTS**
LDAP server with specified FQDN already exists
- **SSL_CERTIFICATE_CHAIN_EMPTY**
No certificates found in input.
- **SSL_CERTIFICATE_HAS_EXPIRED**
SSL certificate has expired.
- **SSL_CERTIFICATE_INVALID_FORMAT**
SSL certificate format is invalid or corrupted.
- **SSL_CERTIFICATE_ISSUER_NOT_FOUND**
SSL certificate issuer not found in certificate chain.
- **SSL_CERTIFICATE_NOT_YET_VALID**
SSL certificate is not yet valid.
- **SSL_CERTIFICATE_VERIFICATION_FAILED**
SSL certificate chain verification failed.
- **SSL_CERTIFICATE_VERIFICATION_INTERNAL_ERROR**
SSL certificate verification has failed because of 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:

```
xcli.py 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 your 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*' has expired on *Expiration Date*.
- USER_IS_PREDEFINED_IN_THE_SYSTEM
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
LDAP role for user is not recognized in the system
- LDAP_SERVER_NOT_FOUND
LDAP server with specified FQDN is not 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_attrb
storage_admin_role
read_only_role
session_cache_period   20
bind_time_limit        20
user_id_attrb          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_attrb=LdapAttrb ] [ xiv_group_attrb=LdapAttrb ]
[ storage_admin_role=LdapRole ] [ read_only_role=LdapRole ]
[ security_admin_role=LdapRole ] [ storage_integration_admin_role=LdapRole ]
[ use_ssl=<yes|no> ] [ user_id_attrb=LdapAttrb ] [ session_cache_period=Minutes ]
[ bind_time_limit=Seconds ] [ first_expiration_event=Days ]
[ second_expiration_event=Days ] [ third_expiration_event=Days ] [ version=LdapVersion ]
[ xiv_user=LdapAttrb ] [ xiv_password=LdapAttrb ]
[ 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_attrb	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_attrb	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.	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 (;).	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 (;).	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 (;).	N	none
use_ssl	Boolean	Indicates whether secure LDAP is mandatory.	N	no
user_id_attrb	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)

Name	Type	Description	Mandatory	Default
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
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.

Example:

```
ldap_config_set
base_dn version xiv_group_attr
storage_admin_role read_only_role
session_cache_period bind_time_limit
use_ssl user_id attrib first_expiration_event
second_expiration_event
third_expiration_event
xiv_user
xiv_password
server_type="SUN DIRECTORY"
```

Output:

Command executed 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 your settings.

- LDAP_CONFIG_CHANGE_IS_ILLEGAL_WHEN_AUTHENTICATION_IS_ACTIVE
This LDAP configuration change is invalid when LDAP configuration is active

Troubleshooting: Disable LDAP-based authentication and then change LDAP configuration.

- LDAP_ROLE_ALREADY_USED
LDAP role is already in use in 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
Dates for expiration events must be in ascending order
- LDAP_READ_ONLY_ROLE_HAS_WRONG_NUMBER_OF_PARTS
LDAP read only role contains too many parts

Troubleshooting: The role is divided to parts by ',' the number of parts should be between 0 and 8

- LDAP_ROLE_HAS_DUPLICATED_PARTS
LDAP role contains duplicated parts
- LDAP_STORAGE_ADMIN_ROLE_HAS_WRONG_NUMBER_OF_PARTS
LDAP storage admin role contains too many parts

Troubleshooting: The role is divided to parts by ',' the number of parts should be between 0 and 8

- LDAP_SECURITY_ADMIN_ROLE_HAS_WRONG_NUMBER_OF_PARTS
LDAP security admin role contains too many parts

Troubleshooting: The role is divided to parts by ',' the number of parts should be between 0 and 8

- LDAP_STORAGE_INTEGRATION_ADMIN_ROLE_HAS_WRONG_NUMBER_OF_PARTS
LDAP storage integration admin role contains too many parts

Troubleshooting: The role is divided to parts by ',' 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

User Category	Permission
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 your 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 executed 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 to 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 your 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 content of a .pem file, with asterisks (*) instead of newlines. In Windows, drag-and-drop the .pem file from the Windows Explorer to the appropriate location in the XCLI session window; the content will be added automatically.	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

Example:

```
ldap_update_server fqdn=ldap.example.com address=1.2.3.4  
remove_certificate=yes
```

Output:

```
Command executed 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**
LDAP server with specified FQDN is not 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 update
- **SSL_CERTIFICATE_CHAIN_EMPTY**
No certificates found in input.
- **SSL_CERTIFICATE_HAS_EXPIRED**
SSL certificate has expired.
- **SSL_CERTIFICATE_INVALID_FORMAT**
SSL certificate format is invalid or corrupted.

- **SSL_CERTIFICATE_ISSUER_NOT_FOUND**
SSL certificate issuer not found in certificate chain.
- **SSL_CERTIFICATE_NOT_YET_VALID**
SSL certificate is not yet valid.
- **SSL_CERTIFICATE_VERIFICATION_FAILED**
SSL certificate chain verification failed.
- **SSL_CERTIFICATE_VERIFICATION_INTERNAL_ERROR**
SSL certificate verification has failed because of 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.xivldap2.com
```

Output:

```
Command executed 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 LDAP server?

Return codes

- **LDAP_SERVER_NOT_FOUND**
LDAP server with specified FQDN is not 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 which has 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.xivldap2.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: AQUAAAAAAAAUVAAYcKhSnhmt01IPSuAbQQAAA==
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
LDAP server with specified FQDN is not 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
LDAP role for 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
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 12 characters consisting of: a-z. A-Z or 0-9. Password is 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, operationsadmin, securityadmin, readonly and opsadmin.	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.

Two predefined users are set system-wide: Admin and Technician.

Example:

```
user_define user=xiv_user1 password=s0mePassw0rd password_verify=s0mePassw0rd
category=applicationadmin
```

Output:

Command executed 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
User name already exists
- MAX_USERS_REACHED
Maximum number of users already defined
- PASSWORDS_DO_NOT_MATCH
Passwords must be identical
- USER_PHONE_NUMBER_MUST_ACCOMPANY_AREA_CODE
Phone numbers and area code must be defined together
- LDAP_AUTHENTICATION_IS_ACTIVE
Command is not available while LDAP authentication is active
- DOMAIN_DOESNT_EXIST

Domain does not exist.

- NON_EXCLUSIVE_USER_NOT_SECURITY_ADMIN
Only a security admin can add a user to a domain non-exclusively.
- USER_CANNOT_BE_ADDED_TO_A_DOMAIN
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 executed 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
Command is not available while LDAP authentication is active
- USER_NAME_DOES_NOT_EXIST
User name does not exist
- USER_CANNOT_BE_DELETED
User cannot be deleted
- USER_IS_REFERRED_TO_BY_DEST
User is referred to by an event destination and therefore cannot be deleted
- USER_OWNS_RECOVERY_KEY
User owns recovery key and therefore cannot be deleted or renamed
- REMOVAL_WOULD_CREATE_UNRESOLVABLE_REFERENCE_BETWEEN_USER_AND_USERGROUP
The operation would result in a user referring to user group that's not in its domain. First remove the reference explicitly.

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 executed 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
User group name does not exist
- USER_NAME_DOES_NOT_EXIST
User name does not exist
- USER_ALREADY_INCLUDED_IN_ANOTHER_GROUP
User is included in another user group
- USER_GROUP_ALREADY_INCLUDES_USER
User group already includes user
- ONLY_APPLICATION_ADMIN_USERS_CAN_BE_GROUPED
User groups can only contain application administrators
- USER_GROUP_HAS_MAXIMUM_NUMBER_OF_USERS
User group already has the maximum number of users
- LDAP_AUTHENTICATION_IS_ACTIVE
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
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 executed 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`
User group name already exists
- `MAX_USER_GROUPS_REACHED`
Maximum number of user groups already defined
- `LDAP_ROLE_ALREADY_USED`
LDAP role is already in use in LDAP configuration or in a user group
- `DOMAIN_DOESNT_EXIST`
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
<code>user_group</code>	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 executed 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 to user group. Are you sure you want to delete this user group?
- ARE_YOU_SURE_YOU_WANT_TO_DELETE_USER_GROUP
Are you sure you want to delete user group?
- ARE_YOU_SURE_YOU_WANT_TO_DELETE_POPULATED_USER_GROUP
One or more internal users are associated to user group. Are you sure you want to delete this user group?

Return codes

- USER_GROUP_NAME_DOES_NOT_EXIST
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
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 executed 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 user from user group?

Return codes

- USER_GROUP_NAME_DOES_NOT_EXIST
User group name does not exist
- USER_NAME_DOES_NOT_EXIST
User name does not exist
- USER_GROUP_DOES_NOT_INCLUDE_USER
User group does not include user
- LDAP_AUTHENTICATION_IS_ACTIVE
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 executed 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

Return codes

- **USER_GROUP_NAME_DOES_NOT_EXIST**
User group name does not exist
- **USER_GROUP_NAME_ALREADY_EXISTS**
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 executed 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**
User group name does not exist
- **LDAP_ROLE_ALREADY_USED**
LDAP role is already in use in LDAP configuration or in a user group
- **DOMAIN_DOESNT_EXIST**
Domain does not exist.
- **REMOVAL_WOULD_CREATE_UNRESOLVABLE_REFERENCE_BETWEEN_HOST_AND_USERGROUP**
The operation would result in a host referring to user group that's not in its domain. First remove the reference explicitly.
- **REMOVAL_WOULD_CREATE_UNRESOLVABLE_REFERENCE_BETWEEN_CLUSTER_AND_USERGROUP**
The operation would result in a cluster referring to user group that's not in its domain. First remove the reference explicitly.
- **REMOVAL_WOULD_CREATE_UNRESOLVABLE_REFERENCE_BETWEEN_USER_AND_USERGROUP**
The operation would result in a user referring to user group that's not in its domain. First remove the reference explicitly.

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
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 executed 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**
User name does not exist
- **USER_NAME_ALREADY_EXISTS**
User name already exists
- **USER_CANNOT_BE_RENAMED**
User cannot be renamed
- **LDAP_AUTHENTICATION_IS_ACTIVE**
Command is not available while LDAP authentication is active
- **USER_OWNS_RECOVERY_KEY**
User owns recovery key and therefore cannot be deleted or renamed
- **OPERATION_NOT_ALLOWED_ON_DESIGNATED_MSM_USER**
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
password	N/A	New password. Users can only change their own passwords. The password consists of 6-12 characters, comprised of a-z, A-Z and 0-9, and is 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 this 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.

The predefined admin user can change all passwords, excluding the user Technician. The category (role) of a user cannot be changed. The user Technician does not require a phone number or email address. Limitations on password changes are as follows:

- Any user can change his/her own password.
- The predefined admin user can change all passwords, excluding the user Technician.
- Passwords are case sensitive.

Example:

```
user_update user=admin password=Passw0rd password_verify=Passw0rd
```

Output:

```
Command executed 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**
User name does not exist
- **PASSWORDS_DO_NOT_MATCH**
Passwords must be identical
- **USER_PHONE_NUMBER_MUST_ACCOMPANY_AREA_CODE**
Phone numbers and area code must be defined together
- **ADMIN_CAN_NOT_CHANGE_TECHNICIAN_USER**
Administrators cannot change the details of the technician user
- **SMS_DESTINATION_REFERS_TO_USER**
An SMS destination refers to the user and therefore it must have a phone number and an area code
- **EMAIL_DESTINATION_REFERS_TO_USER**
An email destination refers to the user and therefore it must have an email address
- **USER_NOT_ALLOWED_TO_CHANGE_OTHER_USERS**
This user is not allowed to change the details of other users
- **USER_NOT_ALLOWED_TO_HAVE_PHONE_NUMBER**
User not allowed to have phone number
- **USER_NOT_ALLOWED_TO_HAVE_EMAIL_ADDRESS**
User not allowed to have 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**
User cannot be updated while LDAP authentication is active
- **NON_EXCLUSIVE_USER_NOT_SECURITY_ADMIN**
Only a security admin can add a user to a domain non-exclusively.
- **PREDEFINED_USER_CANNOT_BE_RESTRICTED_TO_DOMAIN**
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**
User is referred to by an event destination and therefore cannot be deleted
- **OPERATION_NOT_ALLOWED_ON_DESIGNATED_MSM_USER**
Designated MSM user cannot be renamed and cannot be exclusively associated with a domain.
- **REMOVAL_WOULD_CREATE_UNRESOLVABLE_REFERENCE_BETWEEN_USER_AND_USERGROUP**
The operation would result in a user referring to user group that's not in its domain. First remove the reference explicitly.

Creating a new domain

Use the **domain_create** command to create a domain.

```
domain_create domain=DomainName [ hard_capacity=GB soft_capacity=GB ]
[ max_pools=MaxPools ] [ max_volumes=MaxVolumes ] [ max_cgs=MaxCGs ]
```

```
[ max_mirrors=MaxMirrors ] [ max_dms=MaxDataMigrations ]  
[ perf_class=perfClassName ] [ ldap_id=LdapRole ] [ allow_ssd_caching=<yes|no> ]
```

Parameters

Name	Type	Description	Mandatory	Default
domain	Object name	The name of the domain to be created.	Y	N/A
hard_capacity	Integer	Defines the sum of the hard sizes of all the pools associated with the domain, in gigabytes.	N	0
soft_capacity	Integer	Defines the sum of the soft 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
allow_ssd_caching	Boolean	States whether SSD caching for the domain volumes is allowed.	N	yes

Example:

```
xccli -u -c XIV1 domain_create domain=d1 capacity=1000
```

Output:

Command executed 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?

Return codes

- SOFT_SIZE_SMALLER_THAN_HARD_SIZE
Soft size must be equal or larger than hard size
- DOMAIN_ALREADY_EXISTS
A domain with this name already exists.
- DOMAIN_MAX_REACHED
The maximum number of domain objects was reached.
- PERF_CLASS_BAD_NAME
Performance Class does not exist
- DOMAIN_INVALID_CAPACITY
Only soft or hard capacity were defined.
- DOMAIN_INSUFFICIENT_HARD_CAPACITY
There is not enough hard capacity available for the domain.
- DOMAIN_INSUFFICIENT_SOFT_CAPACITY
There is not enough soft 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 cons 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 migration available for the domain.
- PERF_CLASS_ASSOCIATED_WITH_HOSTS
Performance Class *Performance Class* is already in use by host.

Updating a domain definition

Use the **domain_update** command to update a domain definition.

```
domain_update domain=DomainName [ hard_capacity=GB ] [ soft_capacity=GB ]
[ max_pools=MaxPools ] [ max_volumes=MaxVolumes ] [ max_cgs=MaxCGs ]
[ max_mirrors=MaxMirrors ] [ max_dms=MaxDataMigrations ]
[ perf_class=perfClassName ] [ ldap_id=LdapRole ] [ allow_ssd_caching=<yes|no> ]
```

Parameters

Name	Type	Description	Mandatory	Default
domain	Object name	The name of the domain to be updated.	Y	N/A
hard_capacity	Integer	Defines the sum of the hard sizes of all the pools associated with the domain, in gigabytes.	N	0
soft_capacity	Integer	Defines the sum of the soft 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
allow_ssd_caching	Boolean	States whether SSD caching for the domain volumes is allowed.	N	yes

Example:

```
xcli -u -c XIV1 domain_update domain=d1 capacity=10000 max_pools=5
max_volumes=100
```

Output:

Command executed 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?

Return codes

- DOMAIN_DOESNT_EXIST
Domain does not exist.
- PERF_CLASS_BAD_NAME
Performance Class does not exist

- **DOMAIN_MAX_HARD_CAPACITY_TOO_SMALL**
The domain already has more than the specified number of hard capacity.
- **DOMAIN_MAX_SOFT_CAPACITY_TOO_SMALL**
Domain already has more than the specified number of soft capacity.
- **SOFT_SIZE_SMALLER_THAN_HARD_SIZE**
Soft size must be equal or larger than hard size
- **DOMAIN_MAX_VOLUMES_TOO_SMALL**
Domain already has more than the specified number of volumes.
- **DOMAIN_MAX_MIRRORS_TOO_SMALL**
Domain already has more than the specified number of mirrors.
- **DOMAIN_MAX_DMS_TOO_SMALL**
Domain already has more than the specified number of data migrations.
- **DOMAIN_MAX_CGS_TOO_SMALL**
Domain already has more than the specified number of consistency groups.
- **DOMAIN_MAX_POOLS_TOO_SMALL**
Domain already has more than the specified number of pools.
- **DOMAIN_INVALID_CAPACITY**
Only soft or hard capacity were defined.
- **DOMAIN_INSUFFICIENT_HARD_CAPACITY**
There is not enough hard capacity available for the domain.
- **DOMAIN_INSUFFICIENT_SOFT_CAPACITY**
There is not enough soft 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 cons groups available for the domain.
- **DOMAIN_INSUFFICIENT_DMS**
There are not enough data migration available for the domain.
- **PERF_CLASS_ASSOCIATED_WITH_HOSTS**
Performance Class *Performance Class* is already in use by host.

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

User Category	Permission
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- DOMAIN_ALREADY_EXISTS
A domain with this name already exists.
- DOMAIN_DOESNT_EXIST
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
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

Use the **domain_list** command to 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_STJTAG_Gen3_XCLI_r_command_accesscontrol_domain_list_4e414e00000">
  <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="50713d000000"/>
  <perf_class value="QoS1"/>
  <dn value="Domain1"/>
</domain>
```

Field ID	Field output	Default position
name	Name	1
ldap_id	LDAP ID	2
soft_capacity	Soft	3
soft_capacity_MiB	Soft (MiB)	N/A
free_soft_capacity	Free Soft	4
free_soft_capacity_MiB	Free Soft (MiB)	N/A
used_soft_capacity	Used Soft	5
used_soft_capacity_MiB	Used Soft (MiB)	N/A
hard_capacity	Hard	6
hard_capacity_MiB	Hard (MiB)	N/A
free_hard_capacity	Free Hard	7
free_hard_capacity_MiB	Free Hard (MiB)	N/A
used_hard_capacity	Used Hard	8
used_hard_capacity_MiB	Used Hard (MiB)	N/A
max_pools	Max Pools	9
used_pools	Pools	10
max_volumes	Max Volumes	11
used_volumes	Volumes	12
max_mirrors	Max Mirrors	13
used_mirrors	Mirrors	14
max_dms	Max Data Migrations	15
used_dms	Data Migrations	16
max_cgs	Max CGs	17
used_cgs	CGs	18
perf_class	Performance Class	19
allow_ssd_caching	Allow SSD Caching	20

Field ID	Field output	Default position
managed	Managed	21
max_gps	Max Grouped Pools	N/A
used_gps	Grouped Pools	N/A
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
Technicians	Allowed

Return codes

- DOMAIN_DOESNT_EXIST
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
Domain does not exist.
- TOO_MANY_OBJECTS
There are too many objects to output. Run the command again, this time use 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
soft_capacity	Soft	3
soft_capacity_MiB	Soft (MiB)	N/A
free_soft_capacity	Free Soft	4
free_soft_capacity_MiB	Free Soft (MiB)	N/A
used_soft_capacity	Used Soft	5
used_soft_capacity_MiB	Used Soft (MiB)	N/A
hard_capacity	Hard	6
hard_capacity_MiB	Hard (MiB)	N/A
free_hard_capacity	Free Hard	7
free_hard_capacity_MiB	Free Hard (MiB)	N/A
used_hard_capacity	Used Hard	8

Field ID	Field output	Default position
used_hard_capacity_MiB	Used Hard (MiB)	N/A
max_pools	Max Pools	9
used_pools	Pools	10
max_volumes	Max Volumes	11
used_volumes	Volumes	12
max_mirrors	Max Mirrors	13
used_mirrors	Mirrors	14
max_dms	Max Data Migrations	15
used_dms	Data Migrations	16
max_cgs	Max CGs	17
used_cgs	CGs	18
perf_class	Performance Class	19
allow_ssd_caching	Allow SSD Caching	20
managed	Managed	21
max_gps	Max Grouped Pools	N/A
used_gps	Grouped Pools	N/A
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 executed 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
Domain does not exist.
- TARGET_BAD_NAME
Target name does not exist

- **HOST_BAD_NAME**
Host name does not exist
- **CLUSTER_BAD_NAME**
Cluster name does not exist
- **USER_GROUP_NAME_DOES_NOT_EXIST**
User group name does not exist
- **DEST_NAME_DOES_NOT_EXIST**
Destination name does not exist
- **DESTGROUP_NAME_DOES_NOT_EXIST**
Destination group name does not exist
- **EVENT_RULE_NAME_DOES_NOT_EXIST**
Event rule name does not exist
- **USER_IS_NOT_IN_DESTINATION_DOMAINS**
User must be included in the destination domains.
- **DESTINATION_IS_NOT_IN_DESTGROUP_DOMAINS**
Destinations must be included in the destination group domains.
- **DESTINATION_IS_NOT_IN_RULE_DOMAINS**
Destination must be included in the rule domains.
- **DESTGROUP_IS_NOT_IN_RULE_DOMAINS**
Destination groups must be included in the rule domains.
- **ESCALATION_RULE_NOT_IN_RULE_DOMAINS**
Escalation rule must belong to rule domains.
- **CLUSTER_HAS_HOSTS_UNASSOCIATED_WITH_DOMAIN**
Cannot attach a cluster with hosts that aren't 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 executed 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**

Domain does not exist.

- **RESOURCE_NOT_ASSOCIATED_WITH_THIS_DOMAIN**
The resource is not associated with this domain.
- **DOMAIN_VOLUME_MAPPED_TO_HOST**
Host has a volume in the domain mapped to it.
- **DOMAIN_VOLUME_MAPPED_TO_CLUSTER**
Cluster has a volume in the domain mapped to it.
- **REMOVAL_WOULD_CREATE_UNRESOLVABLE_REFERENCE_BETWEEN_HOST_AND_USERGROUP**
The operation would result in a host referring to user group that's not in its domain. First remove the reference explicitly.
- **REMOVAL_WOULD_CREATE_UNRESOLVABLE_REFERENCE_BETWEEN_CLUSTER_AND_USERGROUP**
The operation would result in a cluster referring to user group that's not in its domain. First remove the reference explicitly.
- **HOST_PART_OF_ATTACHED_CLUSTER**
The host is a part of a cluster and cannot be handled individually.
- **DOMAIN_TARGET_IN_USE**
Attempt to remove a target that is used in a mirror or DM relation with a volume in the domain.
- **DOMAIN_VOLUME_BOUND_TO_HOST**
Host has a volume in the domain bound to it via an ALU.
- **TARGET_BAD_NAME**
Target name does not exist
- **HOST_BAD_NAME**
Host name does not exist
- **CLUSTER_BAD_NAME**
Cluster name does not exist
- **USER_GROUP_NAME_DOES_NOT_EXIST**
User group name does not exist
- **DEST_NAME_DOES_NOT_EXIST**
Destination name does not exist
- **DESTGROUP_NAME_DOES_NOT_EXIST**
Destination group name does not exist
- **EVENT_RULE_NAME_DOES_NOT_EXIST**
Event rule name does not exist
- **DETACH_WOULD_MAKE_OBJECT_INACCESSIBLE**
Detaching the object would leave it unassociated with any domain, hence inaccessible.

Troubleshooting: If it's no longer needed, please delete it.

- **USER_IS_NOT_IN_DESTINATION_DOMAINS**
User must be included in the destination domains.
- **DESTINATION_IS_NOT_IN_RULE_DOMAINS**
Destination must be included in the rule domains.
- **DESTINATION_IS_NOT_IN_DESTGROUP_DOMAINS**
Destinations must be included in the destination group domains.
- **REMOVAL_WOULD_CREATE_UNRESOLVABLE_REFERENCE_BETWEEN_USER_AND_USERGROUP**
The operation would result in a user referring to user group that's not in its domain. First remove the reference explicitly.
- **DESTGROUP_IS_NOT_IN_RULE_DOMAINS**
Destination groups must be included in the rule domains.
- **ESCALATION_RULE_NOT_IN_RULE_DOMAINS**
Escalation rule must belong to rule domains.
- **DOMAIN_SCHEDULE_IN_USE**
Cannot move the schedule to other domain since it is in use.
- **DOMAIN_PROXY_VOLUME_MAPPED_TO_HOST**
Host has a proxy volume in the domain mapped to it.
- **DOMAIN_PROXY_VOLUME_MAPPED_TO_CLUSTER**
Cluster has a proxy volume in the domain mapped to it.

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 executed 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
Command is not available while LDAP authentication is active
- DOMAIN_DOESNT_EXIST
Domain does not exist.
- USER_NAME_DOES_NOT_EXIST
User name does not exist
- NON_EXCLUSIVE_USER_NOT_SECURITY_ADMIN
Only a security admin can add a user to a domain non-exclusively.
- USER_CANNOT_BE_ADDED_TO_A_DOMAIN
Specified user cannot be associated with a domain.
- PREDEFINED_USER_CANNOT_BE_RESTRICTED_TO_DOMAIN
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
User is referred to by an event destination and therefore cannot be deleted
- OPERATION_NOT_ALLOWED_ON_DESIGNATED_MSM_USER
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
The operation would result in a user referring to user group that's not in its domain. First remove the reference explicitly.

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

Name	Type	Description	Mandatory
user	Object name	The name of the user.	Y

Example:

```
domain_remove_user domain=d1 user=d1_admin
```

Output:

```
Command executed 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**
User is only associated with this domain. This will delete the user from the system. Are you sure?

Return codes

- **USER_NAME_DOES_NOT_EXIST**
User name does not exist
- **DOMAIN_DOESNT_HAVE_THE_USER**
User is not attached to this domain.
- **DOMAIN_DOESNT_EXIST**
Domain does not exist.
- **LDAP_AUTHENTICATION_IS_ACTIVE**
Command is not available while LDAP authentication is active
- **REMOVAL_WOULD_CREATE_UNRESOLVABLE_REFERENCE_BETWEEN_USER_AND_USERGROUP**
The operation would result in a user referring to user group that's not in its domain. First remove the reference explicitly.
- **DOMAIN_USER_CANNOT_REMOVE_HIMSELF**
Users cannot remove themselves from a domain.
- **USER_IS_REFERRED_TO_BY_DEST**
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 executed 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

Return codes

- **POOL_DOES_NOT_EXIST**
Storage Pool does not exist
- **NO_FREE_HARD_CAPACITY_IN_DOMAIN**
There is not enough free hard space in the domain.
- **NO_FREE_SOFT_CAPACITY_IN_DOMAIN**
There is not enough free soft space in the domain.
- **DOMAIN_DOESNT_EXIST**
Domain does not exist.
- **POOL_ALREADY_ASSOCIATED_WITH_A_DOMAIN**
The pool is already associated with a domain.
- **DOMAIN_MAX_POOLS_REACHED**
The maximum number of domain pools was 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 that is 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.
- **POOL_BELONGS_TO_A_GROUPED_POOL**
Pool belongs to a Grouped Pool.
- **BOUND_ALUS_NOT_IN_DESTINATION**
An ALU that is bound to a volume in the pool is not associated with the target domain.

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 executed 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**
Domain does not exist.
- **POOL_DOES_NOT_EXIST**
Storage Pool does not exist
- **DOMAIN_DOESNT_HAVE_THE_POOL**
Pool is not attached to this domain.
- **MAX_POOLS_REACHED**
Maximum number of Storage Pools already defined
- **MAX_MIRRORS_REACHED**
Maximum number of mirrors already defined
- **MAX_CONS_GROUPS_REACHED**
Maximum number of Consistency Groups already defined.
- **MAX_VOLUMES_REACHED**
Maximum number of volumes already defined
- **MAX_DMS_REACHED**
Maximum number of remote volumes (mirror/migration) is already defined

Troubleshooting: Delete unnecessary Data Migration objects

- **POOL_BELONGS_TO_A_GROUPED_POOL**
Pool belongs to a Grouped Pool.
- **NO_SOFT_SPACE**
The system does not have enough free soft space for the requested Storage Pool soft size
- **NO_HARD_SPACE**
The system does not have enough free hard space for the requested Storage Pool hard size
- **DOMAIN_USED_SCHEDULE_NOT_IN_DESTINATION**
A schedule that is 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 executed 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**
Source domain does not exist.
- **DESTINATION_DOMAIN_DOES_NOT_EXIST**
Destination domain does not exist.
- **POOL_DOES_NOT_EXIST**
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 number of domain pools was 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.
- **POOL_BELONGS_TO_A_GROUPED_POOL**
Pool belongs to a Grouped Pool.
- **NO_FREE_HARD_CAPACITY_IN_DOMAIN**
There is not enough free hard space in the domain.
- **NO_FREE_SOFT_CAPACITY_IN_DOMAIN**
There is not enough free soft space in the domain.
- **DOMAIN_USED_SCHEDULE_NOT_IN_DESTINATION**
A schedule that is 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.
- **BOUND_ALUS_NOT_IN_DESTINATION**
An ALU that is bound 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.
- **SOURCE_AND_DESTINATION_DOMAINS_MUST_BE_DIFFERENT**
Source and destination domains must be different.
- **DOMAIN_CONTAINS_OLVM_VOLUME**
The domain contains a volume in olvm process.

Moving a grouped pool from one domain to another

Use the **domain_move_gp** to move a pool from one domain to another.

```
domain_move_gp gp=gpName src_domain=DomainName dst_domain=DomainName [ adjust=<yes|no> ]
```

Parameters:

Name	Type	Description	Mandatory	Default
gp	Object name	The name of the grouped pool.	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	Adjusts domain resources. If set to <i>true</i> , the domain resources are adjusted to accommodate the pool being moved.	N	no

Example:

```
domain_move_pool pool=p1 src_domain=d1 dst_domain=d2
```

Output:

```
Command executed 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

- SOURCE_DOMAIN_DOES_NOT_EXIST
Source domain does not exist.
- DESTINATION_DOMAIN_DOES_NOT_EXIST
Destination domain does not exist.
- GROUPED_POOL_DOES_NOT_EXIST
Grouped Pool does not exist.
- GROUPED_POOL_NOT_ASSOCIATED_WITH_SOURCE_DOMAIN
The grouped pool is not associated with the source domain.
- DOMAIN_MAX_POOLS_REACHED
The maximum number of domain pools was 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_HARD_CAPACITY_IN_DOMAIN
There is not enough free hard space in the domain.
- NO_FREE_SOFT_CAPACITY_IN_DOMAIN
There is not enough free soft space in the domain.
- DOMAIN_USED_SCHEDULE_NOT_IN_DESTINATION
A schedule that is 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_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_MAX_GROUPED_POOLS_REACHED
Domain exceed the max number of grouped pools allowed.
- 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.

- **NO_GROUPED_POOL_ALLOWED_IN_DOMAIN**
The domain definition does not allow grouped pools.
- **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.
- **BOUND_ALUS_NOT_IN_DESTINATION**
An ALU that is bound to a volume in the pool is not associated with the target domain.
- **SOURCE_AND_DESTINATION_DOMAINS_MUST_BE_DIFFERENT**
Source and destination domains must be different.

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_manage domain=d1 managed=yes
```

Output:

```
Command executed 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**
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=closed
```

Output:

```
Command executed 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

- UNRECOGNIZED_CONFIG_PARAMETER
Unrecognized configuration parameter: 'name'.

Troubleshooting: Use a valid configuration parameter as an input.

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

User Category	Permission
Technicians	Allowed

Return codes

- UNRECOGNIZED_CONFIG_PARAMETER
Unrecognized configuration parameter: 'name'.
Troubleshooting: Use a valid configuration parameter as an input.
- CONF_SERVER_UNREACHABLE
Configuration server 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 executed 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
User name does not exist
- USER_IS_ONLY_DOMAIN_ADMIN
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

User Category	Permission
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

The advanced set also allows application administrators to run the following 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

User Category	Permission
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 connectivity to hosts](#)
Use the **host_connectivity_list** command to list FC and iSCSI-level connectivity to a pre-defined host.

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 **role=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	Disallowed

Return codes

- COMPONENT_IS_NOT_AN_FC_PORT
Component must specify an FC port
- FC_PORT_DOES_NOT_EXIST
Specified FC port does not exist

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
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
port_up_cnt	Port Up	N/A
loss_sync_cnt	Loss Sync	N/A
loss_signal_cnt	Loss Signal	N/A
PER_cnt	PER	N/A

Field ID	Field output	Default position
BER_cnt	BER	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 connectivity to hosts

Use the **host_connectivity_list** command to list FC and iSCSI-level connectivity to a pre-defined host.

```
host_connectivity_list [ host=HostName | fc_host_port=WWPN ]
[ module=ModuleNumber | fcport=ComponentId ] [ domain=DomainName ]
```

Parameters

Name	Type	Description	Mandatory	Default
host	Object name	Limits viewing to the ports of a specific host.	N	All hosts.
fc_host_port	N/A	Limits viewing to this specific port.	N	All ports
module	N/A	Limits output only to the enabled connectivity to this module.	N	All modules
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 an FC 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

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Hardware maintenance commands

This section describes the command-line interface (CLI) for maintaining hardware components.

- [Displaying an ATS configuration](#)
Use the **ats_list** to display information about the components of an Automatic Transfer Switch (ATS).
- [Displaying the system's coupling facilities](#)
Use the **cf_list** to display the list of coupling facilities (CF) in the storage system.
- [Listing system components](#)
Use the **component_list** command to list system components and their status.
- [Listing system components that require service](#)
Use the **component_service_required_list** command to list system components and their status.
- [Listing InfiniBand host card adapters in the storage system](#)
Use the **hca_list** command to list the InfiniBand host card adapters (HCAs) in the storage system's modules.
- [Listing CNA adapters in the system](#)
Use the **cna_list** command to list CNA adapters in the storage system.
- [Listing DIMMs in the system](#)
Use the **dimmm_list** command to list the DIMMs in the storage system.
- [Listing CPUs in the system](#)
Use the **cpu_list** command to list the central processing units (CPU) in the storage system.
- [Listing maintenance modules in the system](#)
Use the **mm_list** command to list maintenance modules in the storage system.
- [Listing NICs in the system](#)
Use the **nic_list** command to list the NICs in the storage system.
- [Listing the internal temperature of modules](#)
Use the **module_temperature_list** command to list the modules' internal temperatures in the storage system.
- [Monitoring the rebuilding or redistribution processes](#)
Use the **monitor_redist** command to monitor the status of the rebuilding or redistribution process.
- [Listing disk status](#)
Use the **disk_list** command to list special disk statuses.
- [Listing module configuration](#)
Use the **module_list** command to list the configuration of all or specified modules.
- [Checking serial connections of modules](#)
Use the **serial_console_check** to check serial connections between modules.
- [Listing serial console statuses](#)
Use the **serial_consoles_list** command to list serial consoles.
- [Displaying the status of a UPS component](#)
Use the **ups_list** command to display the status of a UPS component.
- [Listing service statuses](#)
Use the **service_list** command to list all service specific statuses.
- [Listing PSUs in the system](#)
Use the **psu_list** command to list the PSUs in the system.
- [Creating a trace snapshot](#)
Use the **traces_snapshot** command to create a trace snapshot.
- [Listing trace snapshot on a module](#)
Use the **traces_snapshot_list** command to list trace snapshots on a module.
- [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 fans in the system](#)
Use the **fan_list** command to list fans in the system.
- [Listing SSDs used as flash cache](#)
Use the **ssd_list** command to list SSDs used as flash cache in the storage system.
- [Getting the default state of the SSD caching](#)
Use the **vol_default_ssd_caching_get** command to get the default state of the SSD caching.
- [Setting the default for SSD caching](#)
Use the **vol_default_ssd_caching_set** command to set the default state for SSD caching.
- [Overriding the SSD caching state](#)
Use the **vol_ssd_caching_set** command to override the default SSD caching state for a volume.
- [Displaying the system's average power consumption](#)
Use the **system_average_power_consumption** command to display the storage system's average power consumption.
- [Displaying the system's average temperature](#)
Use the **system_average_temperature** command to display the system's average temperature.

Displaying an ATS configuration

Use the **ats_list** to display information about the components of an Automatic Transfer Switch (ATS).

```
ats_list [ ats=ComponentId ]
```

Parameters:

Name	Description	Mandatory	Default

Name	Description	Mandatory	Default
ats	Lists the configuration of the specified ATS.	N	All ATSs

Example:

```
ats_list ats
```

Output:

Component ID	Status	Currently Functioning	Model	L1 Input OK	L2 Input OK
1:ATS:1	OK	yes	ATS-60A	no	yes

Cont.:

Outlet 1 State	Outlet 2 State	Outlet 3 State	Firmware Version
J2	J2	J2	4

Cont.:

3-Phase	Dual Active
no	no

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
model	Model	4
input_l1	L1 Input OK	5
input_l2	L2 Input OK	6
out1_state	Outlet 1 State	7
out2_state	Outlet 2 State	8
out3_state	Outlet 3 State	9
firmware_version	Firmware Version	10
three_phase_type	3-Phase	11
dual_active_type	Dual Active	12
coil_l1	Coil A On	N/A
coil_l2	Coil B On	N/A
pickup_l1	A Pick-Up	N/A
pickup_l2	B Pick-Up	N/A
default_cal	Default Calibration	N/A
serial_control	Serial Control	N/A
logic_power	Logic Power	N/A
output_30A_1	Output 30A #1	N/A
output_30A_2	Output 30A #2	N/A
output_30A_3	Output 30A #3	N/A
output_10A	Output 10A	N/A
ats_reply_errors	ATS Reply Errors	N/A
ats_connect_errors	ATS Connect Errors	N/A
requires_service	Requires Service	N/A
service_reason	Service Reason	N/A
us_type	US Type	N/A
l1_source	J1 Source	N/A
l2_source	J2 Source	N/A
interlock_failed	Interlock Failed	N/A
p1_current_fault	P1 Current Fault	N/A
p2_current_fault	P2 Current Fault	N/A
p3_current_fault	P3 Current Fault	N/A
coil_c	Coil C On	N/A
coil_d	Coil D On	N/A
pickup_c	C Pick-Up	N/A
pickup_d	D Pick-Up	N/A
oc_j1_a	Over-Current J1 Phase A	N/A
oc_j1_b	Over-Current J1 Phase B	N/A
oc_j1_c	Over-Current J1 Phase C	N/A
oc_j2_a	Over-Current J2 Phase A	N/A
oc_j2_b	Over-Current J2 Phase B	N/A
oc_j2_c	Over-Current J2 Phase C	N/A
no_oc_switching	No OC Switching	N/A
rms_current_p1	RMS Current Outlet P1	N/A
rms_current_p2	RMS Current Outlet P2	N/A

Field ID	Field output	Default position
rms_current_p3	RMS Current Outlet P3	N/A
firmware_l1_version	Firmware J1 Version	N/A
firmware_l2_version	Firmware J2 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

Displaying the system's coupling facilities

Use the **cf_list** to display the list of coupling facilities (CF) in the storage system.

```
cf_list [ module=ModuleNumber | cf=ComponentId ]
```

Parameters:

Name	Description	Mandatory	Default
module	Limits the listing to a specific module.	N	All CFs in all modules.
cf	Lists only a specific CF.	N	A specific CF.

Example:

```
cf_list -f all
```

Output:

Component ID	Status	Currently Functioning	Hardware Status	Serial	Part #
1:CF:10:1	OK	yes	OK	0_521134A5	TRANSCEND_20070418
1:CF:11:1	OK	yes	OK	0_5211349C	TRANSCEND_20070418
1:CF:12:1	OK	yes	OK	0_521133F1	TRANSCEND_20070418
1:CF:13:1	OK	yes	OK	0_521133DF	TRANSCEND_20070418
1:CF:14:1	OK	yes	OK	0_52113389	TRANSCEND_20070418
1:CF:15:1	OK	yes	OK	0_521134AE	TRANSCEND_20070418
1:CF:1:1	OK	yes	OK	0_5211347A	TRANSCEND_20070418
1:CF:2:1	OK	yes	OK	0_521133C0	TRANSCEND_20070418
1:CF:3:1	OK	yes	OK	0_521133B0	TRANSCEND_20070418
1:CF:4:1	OK	yes	OK	0_52113568	TRANSCEND_20070418
1:CF:5:1	OK	yes	OK	0_5211357D	TRANSCEND_20070418
1:CF:6:1	OK	yes	OK	0_5211330F	TRANSCEND_20070418
1:CF:7:1	OK	yes	OK	0_521133D6	TRANSCEND_20070418
1:CF:8:1	OK	yes	OK	0_52113C99	TRANSCEND_20070418
1:CF:9:1	OK	yes	OK	0_5211344C	TRANSCEND_20070418

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
size	Size	N/A
requires_service	Requires Service	N/A
service_reason	Service Reason	N/A

Access control

User Category	Permission
Storage administrator	Allowed

User Category	Permission
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
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

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
```

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 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`.

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
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 InfiniBand host card adapters in the storage system

Use the **hca_list** command to list the InfiniBand host card adapters (HCAs) in the storage system's modules.

```
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 Board ID Part Number
-----
1:HCA:10:1    OK      yes          FALCON DDR      MT_0D80120008  MHRH29B-XTR
1:HCA:13:1    OK      yes          OWL DDR         MT_0F90120008  MHRH2A-XTR
1:HCA:1:1     OK      yes          FALCON DDR      MT_0D80120008  MHRH29B-XTR
1:HCA:4:1     OK      yes          FALCON DDR      MT_0D80120008  MHRH29B-XTR
1:HCA:7:1     OK      yes          FALCON DDR      MT_0D80120008  MHRH29B-XTR

```

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

Field ID	Field output	Default position
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	Board ID	Part Number
1:CNA:4:1	OK	yes	ConnectX-2 SFP+	IBM1080000010	81Y9992
1:CNA:5:1	OK	yes	ConnectX-2 SFP+	IBM1080000010	81Y9992
1:CNA:6:1	OK	yes	Hawk Dual Port	IBM1080000010	MNPH29D-XTR
1:CNA:7:1	OK	yes	ConnectX-2 SFP+	IBM1080000010	81Y9992
1:CNA:8:1	OK	yes	ConnectX-2 SFP+	IBM1080000010	81Y9992
1:CNA:9:1	OK	yes	ConnectX-2 SFP+	IBM1080000010	81Y9992

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

Field ID	Field output	Default position
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 DIMMs in the system

Use the **dimmm_list** command to list the DIMMs in the storage system.

```
dimmm_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:

```
dimmm_list -f all
```

Output:

Component ID	Status	Currently Functioning	Hardware Status	DIMM Id	CPU
1:DIMM:10:1	OK	yes	OK	2	0
1:DIMM:10:2	OK	yes	OK	4	0
1:DIMM:10:3	OK	yes	OK	6	0
1:DIMM:11:1	OK	yes	OK	2	0
1:DIMM:11:2	OK	yes	OK	4	0
1:DIMM:11:3	OK	yes	OK	6	0

Cont.:

Channel	Slot	Size (Mb)	Speed (MHz)	Manufacturer	Serial	Original Serial
0	0	8192	1333	Samsung	50F84144	50F84144
1	0	8192	1333	Samsung	59F84144	59F84144
2	0	8192	1333	Samsung	FAF74144	FAF74144
0	0	8192	1333	Samsung	4AF84144	4AF84144
1	0	8192	1333	Samsung	A1F74144	A1F74144
2	0	8192	1333	Samsung	90F74144	90F74144

Cont.:

Part #	Original Part Number	Requires Service
M393B1K70CH0-YH9	M393B1K70CH0-YH9	
M393B1K70CH0-YH9	M393B1K70CH0-YH9	
M393B1K70CH0-YH9	M393B1K70CH0-YH9	
M393B1K70CH0-YH9	M393B1K70CH0-YH9	
M393B1K70CH0-YH9	M393B1K70CH0-YH9	
M393B1K70CH0-YH9	M393B1K70CH0-YH9	

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
channel	Channel	N/A

Field ID	Field output	Default position
slot	Slot	N/A
size	Size(Mb)	N/A
speed	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 system

Use the **cpu_list** command to list the central processing units (CPU) in the storage system.

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
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

Field ID	Field output	Default position
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 maintenance modules in the system

Use the **mm_list** command to list maintenance modules in the storage system.

```
mm_list [ mm=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
mm	Lists only a specific mm.	N	A specific mm.

Example:

```
mm_list -f all
```

Output:

Component ID	Status	Currently Functioning	Enabled	Version
1:MaintenanceModule:1	OK	yes	yes	MGMT-4.5

Temperature	Serial	Original Serial	Part #	Original Part Number
49	0123456789	0123456789	0123456789	0123456789

Total Memory	Free Memory	Free disk (/)	Free disk (/var)	Link#1
932172	602096	39031456	201873624	yes

```
Link#2    Requires Service
-----
yes      None
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
enabled	Enabled	N/A
version	Version	N/A

Field ID	Field output	Default position
temperature	Temperature	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
total_memory	Total Memory	N/A
free_memory	Free Memory	N/A
free_disk_root	Free disk (/)	N/A
free_disk_var	Free disk (/var)	N/A
link_1	Link#1	N/A
link_2	Link#2	N/A
requires_service	Requires Service	N/A
service_reason	Service Reason	N/A
last_heartbeat_time	Last Heartbeat Time	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:

```
root@extra-6013742-module-9:~# xcli.py nic_list -f all
root@extra-6013742-module-9:~# xcli.py nic_list -Ud -f all
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 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.

Example:

```
module_temperature_list -f all
```

Output:

```
root@nextra-6002207-module-1:~# xcli.py module_temperature_list
```

Module	Normalized Ambient	Raw Ambient	Midplane	EM Card	Fan Controller
1:Module:1	18	17	20	23	24
1:Module:10	17	22	20	25	20
1:Module:13	18	24	21	26	22
1:Module:4	17	18	20	23	22
1:Module:7	18	20	20	25	26

Cont.:

CPU 1	DIMM 2	DIMM 4	DIMM 6	PCIe	InfiniBand HCA	Fibre Channel Port fc-0
36	29	29	28	31	62	
39	31	31	29	35	65	
40	31	31	30	35	69	
33	30	29	28	34	62	52
36	30	29	28	36	71	50

Cont.:

Fibre Channel Port fc-2

```
-----
```

51

51

Field ID	Field output	Default position
component_id	Module	1
normalized_temperature	Normalized Ambient	2
ses_temperatures.0	Raw Ambient	3
ses_temperatures.1	Midplane	4
ses_temperatures.2	EM Card	5
ses_temperatures.3	Fan Controller	6
ses_temperatures.4	CPU 1	7
ses_temperatures.5	CPU 2	N/A
ses_temperatures.6	DIMM 1	N/A
ses_temperatures.7	DIMM 2	8
ses_temperatures.8	DIMM 3	N/A
ses_temperatures.9	DIMM 4	9
ses_temperatures.10	DIMM 5	N/A
ses_temperatures.11	DIMM 6	10
ses_temperatures.12	DIMM 7	N/A
ses_temperatures.13	DIMM 8	N/A
ses_temperatures.14	DIMM 9	N/A
ses_temperatures.15	DIMM 10	N/A
ses_temperatures.16	DIMM 11	N/A
ses_temperatures.17	DIMM 12	N/A
ses_temperatures.18	PCIe	11
ib_hca_temperature	InfiniBand HCA	12
cna_temperature	CNA	13
sas_controller_temperature	SAS Controller	N/A
fc_adapter_temperature.0	Fibre Channel Port fc-0	14
fc_adapter_temperature.1	Fibre Channel Port fc-1	N/A
fc_adapter_temperature.2	Fibre Channel Port fc-2	15
fc_adapter_temperature.3	Fibre Channel Port fc-3	N/A

Access control

User Category	Permission
Storage administrator	Allowed

User Category	Permission
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

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 rebuild or 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
initial_capacity_to_copy	Initial Capacity to Copy (GB)	2
capacity_remaining_to_copy	Capacity Remaining to Copy (GB)	3
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

Listing disk status

Use the **disk_list** command to list special disk statuses.

disk_list [module=ModuleNumber | disk=ComponentId]

Parameters

Name	Description	Mandatory	Default
module	Limits the listing to a specific module.	N	All disks in all modules.
disk	Disk for which special statuses are to be listed.	N	All disks.

This command lists the statuses of the disk, including:

- Component generic status
- Disk capacity
- Model
- Serial

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
capacity	Capacity	4
target_status	Target Status	5
vendor	Vendor	6
model	Model	7
size	Size	8
serial	Serial	9
firmware	Firmware	10
part_number	Fru	11
group	Group	12
temperature	Temperature	13
encryption_state	Encryption State	14
original_vendor	Original Vendor	N/A
original_model	Original Model	N/A
original_serial	Original Serial	N/A
original_reported_serial	Original Reported Serial	N/A
original_device_identifier	Original Identifier	N/A
original_part_number	Original Fru	N/A
original_firmware	Original Firmware	N/A
original_group	Original Group	N/A
requires_service	Requires Service	N/A
service_reason	Service Reason	N/A
revision	Revision	N/A
original_revision	Original Revision	N/A
drive_pn	Drive P/N	N/A
original_drive_pn	Original Drive P/N	N/A
desc.bgd_scan	Background Scan	N/A
desc.disk_id	Disk 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	Disk Temperature	N/A
security_state	Security State	N/A
security_state_last	Last Security State	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 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

This command lists the following information for each module:

- Generic component status
- Module type
- Number of disks
- 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

Example:

`module_list`

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
disk_bay_count	Data Disks	6
fc_port_count	FC Ports	7
ethernet_port_count	iSCSI Ports	8
normalized_temperature	Temperature	9
serial	Serial	N/A
original_serial	Original Serial	N/A
part_number	Part Number	N/A
original_part_number	Original Part Number	N/A
usm_version	USM	N/A
bmc_version	BMC	N/A
bios_version	BIOS	N/A
fpga_version	FPGA	N/A
ses_version	SES	N/A
pdb_firmware	PDB	N/A
pcm_1_firmware	PSU-1	N/A
pcm_2_firmware	PSU-2	N/A
fan_controller_firmware	Fan Controller	N/A
battery_firmware	Battery Firmware	N/A
sas_version	SAS	N/A
infiniband_hca_version	InfiniBand HCA	N/A
cna_version	CNA	N/A
requires_service	Requires Service	N/A
service_reason	Service Reason	N/A
memory_gb	Mem	N/A
temperature	SES Temperature	N/A
chassis_serial	Chassis Serial	N/A
chassis_part_number	Chassis Part Number	N/A
electronics_serial	Electronics Serial	N/A
electronics_part_number	Electronics Part Number	N/A
module_11s_number	11S Number	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

Checking serial connections of modules

Use the **serial_console_check** to check serial connections between modules.

```
serial_console_check
```

This command checks serial connections between modules and sends events reflecting the connectivity status.

Example:

```
serial_console_check
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Listing serial console statuses

Use the **serial_consoles_list** command to list serial consoles.

```
serial_consoles_list [ monitoring_module_id=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
monitoring_module_id	List the status of the specified module only.	N	All modules

Example:

```
serial_consoles_list
```

Output:

Monitoring Module	Desired Module	Connected Module	Status	Is Module Alive	Is HW Node
1:Module:1	1:Module:3	1:Module:3	GOOD	yes	yes
1:Module:15	1:Module:14	1:Module:14	GOOD	yes	yes
1:Module:2	1:Module:1	1:Module:1	GOOD	yes	yes
1:Module:3	1:Module:2	1:Module:2	NO_SERIAL_DATA	yes	yes
1:Module:4	1:Module:6	1:Module:6	GOOD	yes	yes
1:Module:6	1:Module:5	1:Module:5	MODULE_FAILED	no	no
1:Module:7	1:Module:9	1:Module:9	GOOD	yes	yes
1:Module:8	1:Module:7	1:Module:7	GOOD	yes	yes
1:Module:9	1:Module:8	1:Module:8	GOOD	yes	yes

Field ID	Field output	Default position
monitoring_module_id	Monitoring Module	1
desired_module_id	Desired Module	2
connected_module_id	Connected Module	3
state	Status	4
module_alive	Is Module Alive	5
hw_node_alive	Is HW Node Alive	6

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Displaying the status of a UPS component

Use the **ups_list** command to display the status of a UPS component.

```
ups_list [ ups=ComponentId ]
```

Parameters:

Name	Description	Mandatory	Default
ups	Lists the status only for the specified UPS.	N	All UPS systems.

The following information is provided:

- Generic status
- Input power on: Y/N
- Battery charge level
- Last date of self-test
- Result of last self-test
- Is monitoring enabled
- Last calibration date
- result of last calibration
- Status of UPS
- Date of next self test
- Serial Number
- Load level percent
- Apparent load level percent
- Cable check result

Additional information, available through running **ups_list -f all**, includes:

- Last Calibration Date
- Last Calibration Result
- Next Self Test
- Serial
- Original Serial
- Load % Watts
- Apparent Load % VA
- Minutes Left
- Temperature
- AOS Version
- Self-test Status
- Component Test Status
- Batteries Date
- UPS Manufacture Date
- Cable check result

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
input_power_on	Input Power On	4
runtime_remaining	Runtime Remaining	5
battery_charge_level	Battery Charge Level	6
last_self_test_date	Last Self Test Date	7
last_self_test_result	Last Self Test Result	8
is_enabled	Monitoring Enabled	9
ups_status	UPS Status	10
last_calibration_date	Last Calibration Date	N/A
last_calibration_result	Last Calibration Result	N/A
serial	Serial	N/A
original_serial	Original Serial	N/A
load_power_percent	Load % Watts	N/A
apparent_load_power_percent	Apparent Load % VA	N/A
power_consumption	Power Consumption	N/A
predictive_power_load	Predictive Power Load %	N/A
predictive_remaining_runtime	Predictive Remaining Runtime	N/A
internal_temperature	Temperature	N/A
aos_version	AOS Version	N/A
application_version	Application Version	N/A
firmware_version	Firmware Version	N/A
self_test_status	Self-Test Status	N/A
component_test_status	Component Test Status	N/A
battery_year.0	First Battery Year Born	N/A
battery_week.0	First Battery Week Born	N/A
battery_year.1	Second Battery Year Born	N/A

Field ID	Field output	Default position
battery_week.1	Second Battery Week Born	N/A
battery_serial	First Battery Serial	N/A
original_battery_serial	Original First Battery Serial	N/A
second_battery_serial	Second Battery Serial	N/A
original_second_battery_serial	Original Second Battery Serial	N/A
manufacture_date	UPS Manufacture Date	N/A
requires_service	Requires Service	N/A
service_reason	Service Reason	N/A
cable_check_result	Cable Check Result	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 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

User Category	Permission
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing PSUs in the system

Use the **psu_list** command to list the PSUs in the system.

```
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:

```
psu_list
```

Output:

Component ID	Status	Currently Functioning	Hardware Status	Power
1:PSU:13:1	Failed	no	No Input Power	<N/A>
1:PSU:13:2	OK	yes	OK	92.69W
1:PSU:14:1	Failed	no	No Input Power	<N/A>
1:PSU:14:2	OK	yes	OK	114.13W
1:PSU:15:1	Failed	no	No Input Power	<N/A>
1:PSU:15:2	OK	yes	OK	105.11W
1:PSU:6:1	Failed	no	No Input Power	<N/A>
1:PSU:6:2	OK	yes	OK	140.80W
1:PSU:9:1	Failed	no	No Input Power	<N/A>
1:PSU:9:2	OK	yes	OK	111.33W

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
hardware_status	Hardware Status	4
requires_service	Requires Service	N/A
service_reason	Service Reason	N/A
serial	Serial Number	N/A
original_serial	Original Serial Number	N/A
part_number	Part Number	N/A
original_part_number	Original Part Number	N/A
firmware_version	Firmware	N/A
voltage_12v	Voltage (12V)	N/A
voltage_5v	Voltage (5V)	N/A
current	Current	N/A
power	Power	5
fru_descriptor	FRU descriptor	N/A
original_fru_descriptor	Original FRU descriptor	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

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:

```
xcli.py 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

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

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=IPAddress 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 laptop port. The parameter supports up to 5 IP addresses, as explained below.	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.

The **from** parameter supports up to 5 IP addresses, separated by a comma: **from=IPaddress1 , IPaddress2 , IPaddress3**. All addresses must be of the same type, either IPv4 or IPv6.

The **technician** option can be used in conjunction with IPv4 or IPv6 addresses. Its position among IP addresses is not restricted.

The **any** option cannot be combined with unicast IP addresses. It can be used only on its own.

Example:

```
xiv_support_enable finish=2012-2-3.16:30 from=1.2.3.4 comment="Some reason"
```

Output:

```
Command executed 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 positive and define time greater then the current time.
- XIV_SUPPORT_WORK_INVALID_FINISH
Finish must be greater then start and the current time.
- XIV_SUPPORT_WORK_INVALID_FROM
From must be valid IPv4 or IPv6 address.
- LIST_WITH_MIXED_IPV6_AND_IPV4_NOT_ALLOWED
IP addresses should be all IPv4 or all IPv6 addresses.
- LIST_WITH_ANY_OPTION_AND_SPECIFIC_IP_ADDRESSES_NOT_ALLOWED
IP addresses should be unicast or any.

Disabling XIV Support access

Use the **xiv_support_disable** command to disable XIV Support access.

```
xiv_support_disable
```

If XIV support is enabled for more than one IP address (see [Enabling XIV Support access](#) for more details), issuing the **xiv_support_disable** command will disable all of them.

Example:

```
xiv_support_disable
```

Output:

```
Command executed 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
1	2	3	4	

Field ID	Field output	Default position
enabled	Enabled	1
from	From	2
start	Start	3
finish	Finish	4
comment	Comment	5

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 fans in the system

Use the **fan_list** command to list fans in the system.

```
fan_list [ module=ModuleNumber | fan=ComponentId ]
```

Example:

```
fan_list
```

Output:

Component ID	Status	Currently Functioning
1:Fan:1:1	OK	yes
1:Fan:1:10	OK	yes

```

1:Fan:1:2      OK      yes
1:Fan:1:3      OK      yes
1:Fan:1:4      OK      yes
1:Fan:1:5      OK      yes
1:Fan:1:6      OK      yes
1:Fan:1:7      OK      yes
1:Fan:1:8      OK      yes
1:Fan:1:9      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 SSDs used as flash cache

Use the **ssd_list** command to list SSDs used as flash cache in the storage system.

```
ssd_list [ ssd=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
ssd	Drive for which status is requested.	N	All SSDs.

Example:

```
ssd_list
```

Output:

```

Component ID  Status  Currently Functioning
-----
1:SSD:1:1     OK      yes
1:SSD:10:1    OK      yes
1:SSD:2:1     OK      yes
1:SSD:3:1     OK      yes
1:SSD:4:1     OK      yes
1:SSD:5:1     OK      yes
1:SSD:6:1     OK      yes
1:SSD:7:1     OK      yes
1:SSD:8:1     OK      yes
1:SSD:9:1     OK      yes

```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
capacity	Capacity	4
target_status	Target Status	5
vendor	Vendor	6
model	Model	7
size	Size	8
serial	Serial	9
firmware	Firmware	12
part_number	Fru	13
group	Group	14
temperature	Temperature	15
encryption_state	Encryption State	16

Field ID	Field output	Default position
original_vendor	Original Vendor	N/A
original_model	Original Model	N/A
original_serial	Original Serial	N/A
original_part_number	Original Fru	N/A
original_firmware	Original Firmware	N/A
original_group	Original group	N/A
requires_service	Requires Service	N/A
service_reason	Service Reason	N/A
revision	Revision	N/A
original_revision	Original Revision	N/A
drive_pn	Drive_pn	N/A
original_drive_pn	Original drive pn	N/A
desc.last_sample_serial	Last Sample Serial	N/A
desc.last_sample_time	Last Sample Time	N/A
desc.power_on_hours	Power On Hours	N/A
desc.block_wear_leveling	Block Wear Leveling	N/A
desc.secure_erase_status	Secure Erase Status	N/A
desc.temperature_status. reported_severity	Reported Severity	N/A
desc.temperature_status. reported_temperature	Reported Temperature	N/A
desc.temperature_status. temperature	SSD Temperature	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

Getting the default state of the SSD caching

Use the **vol_default_ssd_caching_get** command to get the default state of the SSD caching.

```
vol_default_ssd_caching_get
```

This default state of the SSD caching can be overridden by the **vol_ssd_caching_set** command.

Example:

```
vol_default_ssd_caching_get
```

Output:

```
Command executed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Return codes

- SSD_CACHING_NOT_ENABLED
SSD caching was not enabled.

Setting the default for SSD caching

Use the **vol_default_ssd_caching_set** command to set the default state for SSD caching.

```
vol_default_ssd_caching_set default=<DEFAULT|ENABLED|DISABLED>
```

Parameters

Name	Type	Description	Mandatory
default	Enumeration	The SSD caching state that will be the default.	Y

This command sets the default value for SSD caching state. If a volume is assigned for SSD caching, this default value will be applied automatically.

Example:

```
vol_default_ssd_caching_set default
```

Output:

```
Command executed 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

Overriding the SSD caching state

Use the **vol_ssd_caching_set** command to override the default SSD caching state for a volume.

```
vol_ssd_caching_set [ vol=VolName ] state=<enabled|disabled|default>
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	The name of the volume.	N	All volumes.
state	Enumeration	The SSD caching state that overrides the default.	Y	N/A

Example:

```
vol_ssd_caching_set vol state
```

Output:

```
Command executed 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_APPLY_THE_SSD_CACHING_DEFAULT_ON_ALL_VOLUMES
Are you sure you want to have all volumes use the default SSD caching setting?
- ARE_YOU_SURE_YOU_WANT_TO_ENABLE_SSD_CACHING_FOR_ALL_VOLUMES
Are you sure you want to enable SSD caching for all volumes?
- ARE_YOU_SURE_YOU_WANT_TO_DISABLE_SSD_CACHING_FOR_ALL_VOLUMES
Are you sure you want to disable SSD caching for all volumes?

Return codes

- **VOLUME_BAD_NAME**
Volume name does not exist.
- **VOLUME_IS_SNAPSHOT**
Operation is not permitted on snapshots.
- **DOMAIN_IS_NOT_ALLOWED_TO_USE_SSD_CACHING**
Trying to set volume SSD caching state in domain which is not allowed to use SSD caching.

Displaying the system's average power consumption

Use the **system_average_power_consumption** command to display the storage system's average power consumption.

```
system_average_power_consumption
```

Example:

```
system_average_power_consumption
```

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 temperature

Use the **system_average_temperature** command to display the system's average temperature.

```
system_average_temperature
```

Example:

```
system_average_temperature
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
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.
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 <i>resolution_unit</i> .	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_throughput	Read Hit Very large - Throughput	4
read_hit_large_iops	Read Hit Large - IOps	5
read_hit_large_latency	Read Hit Large - Latency	6
read_hit_large_throughput	Read Hit Large - Throughput	7
read_hit_medium_iops	Read Hit Medium - IOps	8
read_hit_medium_latency	Read Hit Medium - Latency	9
read_hit_medium_throughput	Read Hit Medium - Throughput	10
read_hit_small_iops	Read Hit Small - IOps	11
read_hit_small_latency	Read Hit Small - Latency	12
read_hit_small_throughput	Read Hit Small - Throughput	13
read_miss_very_large_iops	Read Miss Very large - IOps	14

Field ID	Field output	Default position
read_miss_very_large_latency	Read Miss Very large - Latency	15
read_miss_very_large_throughput	Read Miss Very large - Throughput	16
read_miss_large_iops	Read Miss Large - IOps	17
read_miss_large_latency	Read Miss Large - Latency	18
read_miss_large_throughput	Read Miss Large - Throughput	19
read_miss_medium_iops	Read Miss Medium - IOps	20
read_miss_medium_latency	Read Miss Medium - Latency	21
read_miss_medium_throughput	Read Miss Medium - Throughput	22
read_miss_small_iops	Read Miss Small - IOps	23
read_miss_small_latency	Read Miss Small - Latency	24
read_miss_small_throughput	Read Miss Small - Throughput	25
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_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_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_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_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_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_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_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_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_throughput	Read Memory-Hit Very large - Throughput	52
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_throughput	Read Memory-Hit Large - Throughput	55
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_throughput	Read Memory-Hit Medium - Throughput	58
read_memory_hit_small_iops	Read Memory-Hit Small - IOps	59
read_memory_hit_small_latency	Read Memory-Hit Small - Latency	60
read_memory_hit_small_throughput	Read Memory-Hit Small - Throughput	61
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

Remote port address is illegal or does not belong to the remote target

- **VOLUME_BAD_NAME**
Volume name does not exist
- **STATS_TOO_MANY_SAMPLES**
Requested number of statistics samples is too high
- **TARGET_BAD_NAME**
Target name does not exist
- **COMPONENT_DOES_NOT_EXIST**
Component does not exist
- **HOST_BAD_NAME**
Host name does not exist
- **HOST_PORT_DOES_NOT_EXIST**
Port ID is not defined
- **IPINTERFACE_DOES_NOT_EXIST**
IP Interface name does not exist
- **PERF_CLASS_BAD_NAME**
Performance Class does not exist
- **COMMAND_AMBIGUOUS**
User belongs to more than one domain. Please specify a domain or a specific object.
- **DOMAIN_DOESNT_EXIST**
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.
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

User Category	Permission
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Return codes

- **VOLUME_BAD_NAME**
Volume name does not exist
- **POOL_DOES_NOT_EXIST**
Storage Pool does not exist
- **BAD_TIME_FORMAT**
Bad time format. Should be YYYY-MM-DD[.HH[:MM[:SS]]]
- **END_BEFORE_START**
End Time should be later than Start Time
- **VOLUME_IS_SNAPSHOT**
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 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

User Category	Permission	Condition
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
Referenced object does not exist.
- MAX_METADATA_OBJECTS_REACHED
Maximal number of metadata objects 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, msgw, 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
Referenced object does not exist.
- METADATA_OBJECT_KEY_NOT_FOUND
The specified metadata object does not exist.

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, msgw, 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

Metadata-service commands

This section describes the command-line interface (CLI) for handling the metadata service.

- [Creating a metadata service database](#)
Use the **metadata_service_db_create** command to create a new database.
- [Deleting a metadata service database](#)
Use the **metadata_service_db_delete** command to delete a given database
- [Listing all metadata service databases](#)
Use the **metadata_service_db_list** command to list all databases.
- [Adding a key to the metadata service](#)
Use the **metadata_service_entry_create** command to create a new entry in a given database and store the new data as its value.
- [Deleting a key from the metadata service](#)
Use the **metadata_service_entry_delete** command to delete an entry in a given database.
- [Listing all keys in a database](#)
Use the **metadata_service_entry_list** command to list all keys in a given database.
- [Retrieving a metadata service key](#)
Use the **metadata_service_entry_retrieve** command to return the value of a given metadata service key.
- [Updating the metadata service key](#)
Use the **metadata_service_entry_update** command to update the value of an existing entry in a given database.
- [Enabling the metadata service](#)
Use the **metadata_service_enable** command to enable the metadata service.

Creating a metadata service database

Use the **metadata_service_db_create** command to create a new database.

```
metadata_service_db_create db_name=DBName
```

Parameters

Name	Type	Description	Mandatory
db_name	Object name	Name of the database to create	Y

Example:

```
metadata_service_db_create db_name
```

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

- PNODE_MKDIR_SYSCALL_FAILED
The mkdir system call failed

Troubleshooting: Contact support

Deleting a metadata service database

Use the **metadata_service_db_delete** command to delete a given database

```
metadata_service_db_delete db_name=DBName
```

Parameters

Name	Type	Description	Mandatory
db_name	Object name	The name of the database to be deleted.	Y

Example:

```
metadata_service_db_delete db_name
```

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

Listing all metadata service databases

Use the **metadata_service_db_list** command to list all databases.

```
metadata_service_db_list
```

Lists all databases

Example:

```
metadata_service_db_list
```

Output:

```
db_name=grouped_pool
```

Access control

User Category	Permission
Storage administrator	Disallowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Adding a key to the metadata service

Use the **metadata_service_entry_create** command to create a new entry in a given database and store the new data as its value.

```
metadata_service_entry_create value=ParamValue db_name=DBName key=Key
```

Parameters

Name	Type	Description	Mandatory
db_name	Object name	The name of the target database.	Y
key	Object name	The name of the key to create.	Y
value	String	Value of the metadata service key	Y

Example:

```
metadata_service_entry_create value db_name key
```

Output:

```
status_str = "Command completed successfully"
return:
    new_token = "5bbbbc7600eab4fb6c580fa07f02654b"
```

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

- PNODE_MKDIR_SYSCALL_FAILED
The mkdir system call failed

Troubleshooting: Contact support

Deleting a key from the metadata service

Use the **metadata_service_entry_delete** command to delete an entry in a given database.

```
metadata_service_entry_delete db_name=DBName key=Key token=Token
```

Parameters

Name	Type	Description	Mandatory
db_name	Object name	The name of the target database.	Y
key	Object name	The name of key to be deleted.	Y
token	N/A	The expected token of the given key.	Y

Example:

```
metadata_service_entry_delete db_name key token
```

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

Listing all keys in a database

Use the **metadata_service_entry_list** command to list all keys in a given database.

```
metadata_service_entry_list db_name=DBName [ from_timestamp=int ]
```

Parameters

Name	Type	Description	Mandatory	Default
db_name	Object name	The name of the target database.	Y	N/A
from_timestamp	Integer	Return entries older than the specified timestamp.	N	0

Example:

```
metadata_service_entry_list db_name=my_db
```

Output:

```
key 1:
      db = "my_db"
      name = "key1"
      timestamp = "1366806659"
      token = "5bbbbc7600eab4fb6c580fa07f02654b"
```

Access control

User Category	Permission
Storage administrator	Disallowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Retrieving a metadata service key

Use the **metadata_service_entry_retrieve** command to return the value of a given metadata service key.

```
metadata_service_entry_retrieve key=Key db_name=DBName
```

Parameters

Name	Type	Description	Mandatory
db_name	Object name	The name of the target database.	Y
key	Object name	The name of the key to be retrieved.	Y

Example:

```
metadata_service_entry_retrieve key db_name
```

Output:


```
returned_value = "value"
token = "7c451ea3d2384e3604ad65de18e53c27"
```

Access control

User Category	Permission
Storage administrator	Disallowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Updating the metadata service key

Use the **metadata_service_entry_update** command to update the value of an existing entry in a given database.

```
metadata_service_entry_update value=ParamValue db_name=DBName key=Key token=Token
```

Parameters

Name	Type	Description	Mandatory
token	N/A	The token of the metadata service key.	Y
db_name	Object name	The name of the target database.	Y
key	Object name	The name of the key to be updated.	Y
value	String	New metadata key value.	Y

Example:

```
metadata_service_entry_update value db_name key token
```

Output:

```
new_token = "ccb563c456d21662f0a7ff038077d9c1"
```

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

- PNODE_MKDIR_SYSCALL_FAILED
The mkdir system call failed

Troubleshooting: Contact support

Enabling the metadata service

Use the **metadata_service_enable** command to enable the metadata service.

```
metadata_service_enable
```

Example:

```
metadata_service_enable
```

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
Operations administrator	Disallowed
Technicians	Disallowed
Host side accelerator client	Disallowed

Return codes

- **NO_SPACE**
The system does not have enough free space for the requested Storage Pool size

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 keyserver](#)
Use the **encrypt_keyserver_define** command to define a keyserver to be used by the system.
- [Removing a keyserver](#)
Use the **encrypt_keyserver_delete** command to remove a keyserver used by the system.
- [Displaying keyserver status](#)
Use the **encrypt_keyserver_list** command to list the keyservers currently defined in the system along with their connectivity status.
- [Obtaining a new master key](#)
Use the **encrypt_keyserver_rekey** command to initiate a rekey against the master keyserver.
- [Renaming a keyserver](#)
Use the **encrypt_keyserver_rename** command to change the name of a defined keyserver.
- [Changing keyserver properties](#)
Use the **encrypt_keyserver_update** command to change a keyserver'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 keysevers, 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.

Disabling encryption

Use the **encrypt_disable** command to disable the data protection feature.

encrypt_disable

This command disables the data protection feature. A prerequisite for this is that no volumes are defined in the system. In addition to disabling the data protection, 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 executed 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: A yes option is required for this command

Return codes

- UNSUPPORTED_HARDWARE
Cannot utilize encryption on unsupported hardware.
Troubleshooting: Contact support to verify encryption status.
- ENCRYPT_NOT_ENABLED
Encryption is not enabled.
Troubleshooting: Check that encryption is enabled and try again 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 statistics volume for disabling encryption.
Troubleshooting: Please contact support.
- CANNOT_CRYPTO_ERASE_DISKS
Cannot crypto-erase disks.
Troubleshooting: Contact support.
- CANNOT_WRITE_TO_KEY_REPOSITORY
Failed writing keys to the key repository.
Troubleshooting: Contact support.
- NO_LIVE_KEYSERVER_GATEWAY_NODE
There is no live key server gateway node on the system.
Troubleshooting: Please restart the key server gateway node and try again.
- NO_MASTER_KEYSERVER_DEFINED
There is no master key server defined on the system.
Troubleshooting: Please 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.
- COMPRESSION_RECOVERY_IN_PROGRESS
Operation cannot be completed because another recovery process is in progress

Enabling encryption

Use the **encrypt_enable** command to enable the data protection feature.

```
encrypt_enable [ recovery_keys=<yes|no> ]
```

Parameters

Name	Type	Description	Mandatory	Default
recovery_keys	Boolean	Defines whether recovery keys are required for encryption activation.	N	yes

This command is entered by a security administrator to enable the data protection feature. In order for this command to complete successfully, all of the following prerequisites must be fulfilled: Current encryption state must be DISABLED (displayed as "Supported" in state_list) One master keyserver configured successfully Recovery keys generated and verified, for and by at least 2 separate security administrators, unless the recovery_keys parameter is set to no.

Example:

```
encrypt_enable recovery_keys=yes -y
```

Output:

```
Command executed 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: A yes option is required for this command

Return codes

- UNSUPPORTED_HARDWARE
Cannot utilize encryption on unsupported hardware.
Troubleshooting: Contact support to verify encryption status.
- INVALID_RECOVERY_KEY_STATE
Recovery key state is inconsistent with the option provided.
Troubleshooting: Check the recovery key state using encrypt_recovery_key_status.
- ENCRYPTION_ALREADY_ENABLED
Encryption has already been enabled.
Troubleshooting: Check the state_list command
- CANNOT_ENROLL_SOME_DISKS
Failed software components prevent enrolling some disks.
Troubleshooting: Contact support.
- CANNOT_ENROLL_SOME_SSDS
Failed SSDs cannot be enrolled.
Troubleshooting: Contact support.
- SYSTEM_IS_REDISTRIBUTING
Operation is not allowed during rebuild or phase-in
- NO_LIVE_KEYSERVER_GATEWAY_NODE
There is no live key server gateway node on the system.
Troubleshooting: Please restart the key server gateway node and try again.
- NO_MASTER_KEYSERVER_DEFINED
There is no master key server defined on the system.
Troubleshooting: Please 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.

- COMPRESSION_RECOVERY_IN_PROGRESS
Operation cannot be completed because another recovery process is in progress

Defining a keyserver

Use the **encrypt_keyserver_define** command to define a keyserver 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 keyserver being added.	Y	N/A
certificate	N/A	The public certificate of the keyserver being added.	Y	N/A
master	Boolean	Defines whether this keyserver is the primary keyserver used for key retrieval.	N	no
ipv4	N/A	The IPv4 address of the keyserver being added. Either one IPv4 and/or one IPv6 must be used.	N	NONE
ipv6	N/A	The IPv6 address of the keyserver being added. Either one IPv4 and/or one IPv6 must be used.	N	NONE
port	Integer	Port used for keyserver communication.	N	5696
keyserver_type	Enumeration	The type of the keyserver to communicate with.	N	TKLM

This command defines a keyserver 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 keyserver (but preferably two, and no more than four) must be defined and accessible in order for encrypt_enable to succeed. Only one of the keyservers may be defined as master.

Example:

```
encrypt_keyserver_define
name=snocone ipv4=snocone.ibm.com ipv6=2002::a5a7
certificate="-----BEGIN CERTIFICATE---*MIICyTCCAbGgAwIBAgIXLSiyd2FPMA0GCSqGSIb3IiEBCwUAMBQxEjAQAgNVBVuTCXNrbG5pdHNv*.....
*erD5HgQHSkfr3FEM+b6EBOUPFIBrys8rKtLRbWvovobq*---END CERTIFICATE-----"
```

Note: To input the certificate as one line, make sure to add asterisks (*) at the beginning and the end of each line.

Output:

Command executed 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 utilize encryption on unsupported hardware.
Troubleshooting: Contact support to verify encryption status.
- ENCRYPTION_TOO_MANY_KEYSERVERS
Too many key servers, cannot add.
Troubleshooting: Delete a key server and try again.
- ENCRYPTION_UNSUPPORTED_KEYSERVER_TYPE
Unsupported key server type.
- ENCRYPTION_KEYSERVER_NAME_EXISTS
Key server name already exists.
Troubleshooting: Check the currently defined key servers
- ENCRYPTION_KEYSERVER_MUST_HAVE_ADDRESS
Key server must have at least one address (IPv4/IPv6).
Troubleshooting: Add ipv4= or ipv6= to the command

- **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
- **LOADED_KEYSERVER_CERTIFICATE_TOO_BIG**
Key server not added as the certificate is too large.
Troubleshooting: Please only use one certificate per PEM file.
- **SSL_CERTIFICATE_HAS_EXPIRED**
SSL certificate has expired.
- **SSL_CERTIFICATE_VERIFICATION_FAILED**
SSL certificate chain verification failed.
- **SSL_CERTIFICATE_INVALID_FORMAT**
SSL certificate format is invalid or corrupted.
- **SSL_CERTIFICATE_NOT_YET_VALID**
SSL certificate is not yet valid.
- **SSL_CERTIFICATE_VERIFICATION_INTERNAL_ERROR**
SSL certificate verification has failed because of internal system error.
- **SSL_CERTIFICATE_ISSUER_NOT_FOUND**
SSL certificate issuer not found in certificate chain.
- **SSL_CERTIFICATE_CHAIN_EMPTY**
No certificates found in input.

Removing a keyserver

Use the **encrypt_keyserver_delete** command to remove a keyserver used by the system.

```
encrypt_keyserver_delete name=Name
```

Parameters

Name	Type	Description	Mandatory
name	String	The name of a defined keyserver.	Y

Example:

```
encrypt_keyserver_delete name=snocone
```

Output:

```
Command executed 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 utilize encryption on unsupported hardware.
Troubleshooting: Contact 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 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 keyserver status

Use the **encrypt_keyserver_list** command to list the keysevers currently defined in the system along with their connectivity status.

```
encrypt_keyserver_list [ check_status=<yes|no> ]
```

Parameters

Name	Type	Description	Mandatory	Default
check_status	Boolean	Defines whether to update the status of keyserver communication paths before displaying them.	N	no

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

Cont. :

Address

9.11.236.1

2002::1

snocone.tucson.ibm.com

2002:90b:e006:238:209:6bff:fe00:a5a7

tklm-sa.ibm.com

Field ID	Field output	Default position
module_id	Module	1
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 utilize encryption on unsupported hardware.

Troubleshooting: Contact support to verify encryption status.

Obtaining a new master key

Use the **encrypt_keyserver_rekey** command to initiate a rekey against the master keyserver.

encrypt_keyserver_rekey

This command initiates a rekeying (getting new cryptographic material) with the master keyserver.

Example:

encrypt_keyserver_rekey

Output:

Command executed 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 utilize encryption on unsupported hardware.
Troubleshooting: Contact support to verify encryption status.
- **ENCRYPT_NOT_ENABLED**
Encryption is not enabled.
Troubleshooting: Check that encryption is enabled and try again the command.
- **CANNOT_GET_XIV_MASTER_KEY**
Problem obtaining XIV master key from the key server.
Troubleshooting: Check that the key server is active and serving keys, and contact support.
- **CANNOT_GET_NEW_KEY_REQUEST**
Error requesting encryption keys from the key server gateway node.
Troubleshooting: Check that the key server is actively serving keys
- **CANNOT_UPDATE_KEY_METADATA**
Cannot update metadata in key repository for new key.
Troubleshooting: Contact support.
- **CANNOT_CANNOT_GENERATE_EXMK_ESKH**
Problem generating EXMK and ESKH.
Troubleshooting: Contact support.
- **CANNOT_WRITE_TO_KEY_REPOSITORY**
Failed writing keys to the key repository.
Troubleshooting: Contact support.
- **CANNOT_COPY_KEYS_IN_KEY_REPOSITORY**
Problem copying current keys to old keys location in the key repository.
Troubleshooting: Contact support.
- **ENCRYPTION_KR_WRITE_FAILED**
Error writing to the key repository.
Troubleshooting: Contact support
- **ENCRYPTION_KR_READ_FAILED**
Error reading the key repository.
Troubleshooting: Contact support
- **NO_LIVE_KEYSERVER_GATEWAY_NODE**
There is no live key server gateway node on the system.
Troubleshooting: Please restart the key server gateway node and try again.
- **NO_MASTER_KEYSERVER_DEFINED**

There is no master key server defined on the system.

Troubleshooting: Please 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.

Renaming a keyserver

Use the `encrypt_keyserver_rename` command to change the name of a defined keyserver.

```
encrypt_keyserver_rename name=Name new_name=Name
```

Parameters

Name	Type	Description	Mandatory
new_name	String	The new name of the keyserver.	Y
name	String	The current name of a defined keyserver.	Y

Example:

```
encrypt_keyserver_rename name=nachos new_name=snocone
```

Output:

```
Command executed 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 utilize encryption on unsupported hardware.
Troubleshooting: Contact support to verify encryption status.
- `ENCRYPTION_UNKNOWN_KEYSERVER`
Unknown key server name.
Troubleshooting: Check the currently defined key servers
- `ENCRYPTION_KEYSERVER_NAME_EXISTS`
Key server name already exists.
Troubleshooting: Check the currently defined key servers

Changing keyserver properties

Use the `encrypt_keyserver_update` command to change a keyserver'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 keyserver to be updated.	Y	N/A
certificate	N/A	The public certificate of the keyserver to be updated.	N	none
master	Enumeration	Indicates whether this keyserver is the master.	N	no
ipv4	N/A	The IPv4 address.	N	none
ipv6	N/A	The IPv6 address.	N	none

Name	Type	Description	Mandatory	Default
port	Integer	Port number for communications.	N	5696

This command is used to update a keyserver's address, port, or certificate.

Example:

```
encrypt_keyserver_update name=nachos master=yes ipv4=10.0.0.1
ipv6=2001::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 utilize encryption on unsupported hardware.
Troubleshooting: Contact 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
- **LOADED_KEYSERVER_CERTIFICATE_TOO_BIG**
key server not added as the certificate is too large.
Troubleshooting: Please only use one certificate per PEM file.
- **SSL_CERTIFICATE_HAS_EXPIRED**
SSL certificate has expired.
- **SSL_CERTIFICATE_VERIFICATION_FAILED**
SSL certificate chain verification failed.
- **SSL_CERTIFICATE_INVALID_FORMAT**
SSL certificate format is invalid or corrupted.
- **SSL_CERTIFICATE_NOT_YET_VALID**
SSL certificate is not yet valid.
- **SSL_CERTIFICATE_VERIFICATION_INTERNAL_ERROR**
SSL certificate verification has failed because of internal system error.
- **SSL_CERTIFICATE_ISSUER_NOT_FOUND**
SSL certificate issuer not found in certificate chain.
- **SSL_CERTIFICATE_CHAIN_EMPTY**
No certificates found in 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 key servers. 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 executed 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 utilize encryption on unsupported hardware.
Troubleshooting: Contact support to verify encryption status.
- **INVALID_RECOVERY_KEY_FRAGMENT**
Recovery key fragment given does not match stored key.
Troubleshooting: Verify that the proper key(share) has been used
- **GENERIC_FAILED**
Generic encryption failure.
Troubleshooting: Contact support.
- **INVALID_RECOVERY_KEY_USER**
User is not a valid recovery key administrator.
Troubleshooting: Check that the user names provided are valid
- **NO_LIVE_KEYSERVER_GATEWAY_NODE**
There is no live key server gateway node on the system.
Troubleshooting: Please restart the key server gateway node and try again.
- **CANNOT_READ_FROM_KEY_REPOSITORY**
Failed reading keys from the key repository.
Troubleshooting: Contact support.
- **RK_FAILED_VERIFY_SLEEP**
Too many failed verify attempts, please wait and try again.
Troubleshooting: Wait a little and try again.
- **ENCRYPTION_KR_WRITE_FAILED**
Error writing to the key repository.
Troubleshooting: Contact support
- **RK_ENTER_SYSTEM_STATE_INVALID**
Command is supported in maintenance mode only.
Troubleshooting: Switch system state to maintenance mode.
- **INVALID_RECOVERY_KEY_STATE**
Recovery key state is inconsistent with the option provided.

Troubleshooting: Check the recovery key state using `encrypt_recovery_key_status`.

- `RECOVERY_KEY_ALREADY_VERIFIED`
The recovery key has already been verified.

Troubleshooting: Check the recovery key state using `encrypt_recovery_key_list`.

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 ]
```

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

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 executed 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 utilize encryption on unsupported hardware.
Troubleshooting: Contact support to verify encryption status.
- `NO_LIVE_KEYSERVER_GATEWAY_NODE`
There is no live key server gateway node on the system.
Troubleshooting: Please restart the key server gateway node and try again.
- `CANNOT_WRITE_TO_KEY_REPOSITORY`
Failed writing keys to the key repository.
Troubleshooting: Contact support.
- `CANNOT_GET_NEW_KEY_REQUEST`
Error requesting encryption keys from the key server gateway node.
Troubleshooting: Check 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: try again the command with at least 2 security administrators
- `ENCRYPTION_KR_WRITE_FAILED`

Error writing to the key repository.

Troubleshooting: Contact support

- ENCRYPTION_ALREADY_ENABLED
Encryption has already been enabled.

Troubleshooting: Check the state_list command

- NO_MASTER_KEYSERVER_DEFINED
There is no master key server defined on the system.

Troubleshooting: Please define a master key server by invoking encrypt_key server_update and try again.

- INVALID_RECOVERY_KEY_STATE
Recovery key state is inconsistent with the option provided.

Troubleshooting: Check the recovery key state using encrypt_recovery_key_status.

- INSUFFICIENT_RK_ADMINS
Number of users must be greater than or equal to the minimal required number.

Troubleshooting: try again the command with at least the minimum number of required users

- CANNOT_GENERATE_KEYS_ON_KEYSERVER_GATEWAY
Failed to generate XMK and hashes on key server gateway node.

Troubleshooting: Contact support.

- ENCRYPTION_KR_READ_FAILED
Error reading the key repository.

Troubleshooting: Contact support

- CANNOT_UPDATE_KEY_METADATA
Cannot update metadata in key repository for new key.

Troubleshooting: Contact 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 executed 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
Technicians	Disallowed

Return codes

- UNSUPPORTED_HARDWARE
Cannot utilize encryption on unsupported hardware.

Troubleshooting: Contact support to verify encryption status.

- INVALID_RECOVERY_KEY_USER

User is not a valid recovery key administrator.

Troubleshooting: Check that the user names provided are valid

- NO_LIVE_KEYSERVER_GATEWAY_NODE
There is no live key server gateway node on the system.

Troubleshooting: Please restart the key server gateway node and try again.

- CANNOT_READ_FROM_KEY_REPOSITORY
Failed reading keys from the key repository.

Troubleshooting: Contact support.

- CANNOT_GET_NEW_KEY_REQUEST
Error requesting encryption keys from the key server gateway node.

Troubleshooting: Check 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
There is no master key server defined on the system.

Troubleshooting: Please define a master key server by invoking `encrypt_key server_update` and try again.

- INVALID_RECOVERY_KEY_STATE
Recovery key state is inconsistent with the option provided.

Troubleshooting: Check the recovery key state using `encrypt_recovery_key_status`.

- RECOVERY_KEY_ALREADY_VERIFIED
The recovery key has already been verified.

Troubleshooting: Check the recovery key state using `encrypt_recovery_key_list`.

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
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- UNSUPPORTED_HARDWARE
Cannot utilize encryption on unsupported hardware.

Troubleshooting: Contact support to verify encryption status.

- NO_LIVE_KEYSERVER_GATEWAY_NODE
There is no live key server gateway node on the system.

Troubleshooting: Please restart the key server gateway node and try again.

- CANNOT_READ_FROM_KEY_REPOSITORY
Failed reading keys from the key repository.

Troubleshooting: Contact support.

- INSUFFICIENT_RK_ADMIN_THRESHOLD
Recovery key creation requires at least two security administrators.

Troubleshooting: try again the command with at least 2 security administrators

- ENCRYPTION_KR_WRITE_FAILED
Error writing to the key repository.

Troubleshooting: Contact support

- NO_MASTER_KEYSERVER_DEFINED
There is no master key server defined on the system.

Troubleshooting: Please define a master key server by invoking encrypt_key server_update and try again.

- INVALID_RECOVERY_KEY_STATE
Recovery key state is inconsistent with the option provided.

Troubleshooting: Check the recovery key state using encrypt_recovery_key_status.

- INSUFFICIENT_RK_ADMINS
Number of users must be greater than or equal to the minimal required number.

Troubleshooting: try again the command with at least the minimum number of required users

- CANNOT_GENERATE_KEYS_ON_KEYSERVER_GATEWAY
Failed to generate XMK and hashes on key server gateway node.

Troubleshooting: Contact 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

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 utilize encryption on unsupported hardware.
Troubleshooting: Contact support to verify encryption status.
- **CANNOT_READ_FROM_KEY_REPOSITORY**
Failed reading keys from the key repository.
Troubleshooting: Contact 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 executed 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 utilize encryption on unsupported hardware.
Troubleshooting: Contact support to verify encryption status.
- **INVALID_RECOVERY_KEY_FRAGMENT**
Recovery key fragment given does not match stored key.

Troubleshooting: Verify that the proper key(share) has been used

- **GENERIC_FAILED**
Generic encryption failure.

Troubleshooting: Contact support.

- **INVALID_RECOVERY_KEY_USER**
User is not a valid recovery key administrator.

Troubleshooting: Check that the user names provided are valid

- **CANNOT_READ_FROM_KEY_REPOSITORY**
Failed reading keys from the key repository.

Troubleshooting: Contact support.

- **RK_FAILED_VERIFY_SLEEP**
Too many failed verify attempts, please wait and try again.

Troubleshooting: Wait a little and try again.

- **ENCRYPTION_KR_WRITE_FAILED**
Error writing to the key repository.

Troubleshooting: Contact support

- **INVALID_RECOVERY_KEY_STATE**
Recovery key state is inconsistent with the option provided.

Troubleshooting: Check the recovery key state using `encrypt_recovery_key_status`.

- **RECOVERY_KEY_ALREADY_VERIFIED**
The recovery key has already been verified.

Troubleshooting: Check the recovery key state using `encrypt_recovery_key_list`.

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

User Category	Permission
Technicians	Allowed

Return codes

- **UNSUPPORTED_HARDWARE**
Cannot utilize encryption on unsupported hardware.
Troubleshooting: Contact support to verify encryption status.
- **CANNOT_READ_FROM_KEY_REPOSITORY**
Failed reading keys from the key repository.
Troubleshooting: Contact 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 executed 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: Check that encryption is enabled and try again the command.
- **RK_ENTER_SYSTEM_STATE_INVALID**
Command is supported in maintenance mode only.
Troubleshooting: Switch system state to maintenance mode.

xmirror commands

This section describes the command-line interface (CLI) for xmirror creation and management.

- [Viewing xmirror status](#)
Use the **xmirror_list** command to list the status and configuration of defined xmirrors.
- [Defining an xmirror relation](#)
Use the **xmirror_define** command to define a new xmirror relation.
- [Activating an xmirror](#)
Use the **xmirror_activate** command to activate all the mirrors within the xmirror.
- [Deactivating an xmirror](#)
Use the **xmirror_deactivate xmirror=XmirrorName** command to deactivate all the mirrors in the xmirror relation.
- [Deleting an xmirror relation](#)
Use the **xmirror_delete** command to delete an existing 3-way mirror relation.
- [Renaming an xmirror relation](#)
Use the **xmirror_rename** command to rename an xmirror.

- [Changing roles in an xmirror relation](#)
Use the **xmirror_change_role** command to change a local xmirror peer's role.
- [Defining a standby mirror in an xmirror relation](#)
Use the **xmirror_register_standby_mirror** command to define a standby mirror in an existing xmirror relation.

Viewing xmirror status

Use the **xmirror_list** command to list the status and configuration of defined xmirrors.

```
xmirror_list [ xmirror=XmirrorName ] [ domain=DomainName ]
```

Parameters

Name	Type	Description	Mandatory	Default
xmirror	Object name	xmirror name.	N	[none]
domain	Object name	The domain name.	N	All Domains.

If no parameters are specified, all xmirror objects are listed. When a name is specified, only the specified xmirror is listed.

The following information is displayed:

Field ID	Field output	Default position
name	Name	1
xmirror_global_id	Xmirror ID	2
local_volume_name	Local volume name	3
xmirror_role	Local Xmirror Role	4
xmirror_state	Xmirror State	5
standby_mirror_state	Standby Mirror	6
master_target_name	Master	7
smaster_target_name	SMaster	8
slave_target_name	Slave	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

Defining an xmirror relation

Use the **xmirror_define** command to define a new xmirror relation.

```
xmirror_define xmirror=XmirrorName vol=VolumeName smaster_target=TargetName  
slave_target=TargetName
```

Parameters

Name	Type	Description	Mandatory
xmirror	Object name	xmirror name.	Y
vol	Object name	Volume name.	Y
smaster_target	Object name	smaster target name.	Y
slave_target	Object name	Slave target name.	Y

Assuming that a synchronous mirror of volume V is defined on systems A (master) and B (smaster), and an asynchronous mirror of volume V is defined on systems A (master) and C (slave), this command (when run on A), creates a named xmirror object which is used to refer to the 3-way relation. The name must not exist on any of the systems. When the command successfully completes, the named xmirror object will be recognized on all systems. Note that systems are referred to via their targets.

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

- **XMIRROR_BAD_NAME**
There is no xmirror with this name
- **TARGET_BAD_NAME**
Target name does not exist
- **VOLUME_BAD_NAME**
Volume name does not exist
- **XMIRROR_SMASTER_TARGET_NOT_CONNECTED**
Xmirror SMaster target is not connected, can not complete operation.
- **XMIRROR_SLAVE_TARGET_NOT_CONNECTED**
Xmirror Slave target is not connected, can not complete operation.
- **XMIRROR_SLAVE_MASTER_MIRROR_IS_MISSING**
Missing or bad master mirror uid provided to slave system!
- **XMIRROR_SMASTER_INVALID_CONFIGURATION**
Xmirror substitute-master mirror configuration is invalid
- **XMIRROR_ALREADY_DEFINED**
An Xmirror is already defined for this volume!
- **XMIRROR_MASTER_INVALID_CONFIGURATION**
Xmirror master mirror configuration is invalid
- **XMIRROR_DEFINE_FAILED_TO_ROLLBACK_MANUAL_CLEANUP_REQUIRED**
An error occurred during xmirror define and the system failed to rollback, Please manually cleanup xmirror objects on all systems
- **XMIRROR_MASTER_INVALID_MIRROR_STATE**
Invalid mirror state in master
- **XMIRROR_INIT_MIRROR_BAD_UID**
Failed to init xmirror on remote system because the mirror uid provided was invalid!
- **XMIRROR_SMASTER_MASTER_MIRROR_IS_MISSING**
Missing or bad master mirror uid provided to smaster system!
- **XMIRROR_SMASTER_INVALID_MIRROR_STATE**
Invalid mirror state in substitute master
- **XMIRROR_INVALID_MASTER_SLAVE_TARGET_CONNECTIVITY**
Cannot define xmirror because master-slave target connectivity at master system is invalid, try again in a few seconds
- **XMIRROR_SLAVE_SMASTER_TARGET_MISMATCH**
Given smaster target does not match defined slave-smaster mirror target
- **XMIRROR_SLAVE_INVALID_CONFIGURATION**
Xmirror slave mirror configuration is invalid
- **XMIRROR_SLAVE_INVALID_MIRROR_STATE**
Invalid mirror state in slave
- **TARGET_BAD_TYPE**
Target machine is not XIV machine
- **XMIRROR_MASTER_SLAVE_TARGET_MISMATCH**
Given smaster target does not match defined master-slave mirror target
- **XMIRROR_SMASTER_MASTER_TARGET_MISMATCH**
Given master target does not match defined smaster-master mirror target
- **XMIRROR_SLAVE_MASTER_TARGET_MISMATCH**
Given master target does not match defined slave-master mirror target
- **XMIRROR_INVALID_MASTER_SMASTER_TARGET_CONNECTIVITY**
Cannot define xmirror because master-smaster target connectivity at master system is invalid
- **XMIRROR_MASTER_SMASTER_TARGET_MISMATCH**
Given smaster target does not match defined master-smaster mirror target
- **XMIRROR_SMASTER_SLAVE_TARGET_MISMATCH**

given slave target does not match defined smaster-slave mirror target

- **VOLUME_HAS_STANDBY_SNAPSHOTS**
Xmirror volume has standby snapshots created by a previous process.
- **XMIRROR_NAME_ALREADY_EXISTS_ON_SLAVE**
Xmirror name already exists on slave system, can not create xmirror
- **XMIRROR_MIRRORING_INCOMPATIBLE_TARGET_VERSION**
Xmirror Mirroring is not supported between the system versions of the specified peers.
- **XMIRROR_NUM_OF_XMIRRORS_ON_SLAVE_LIMIT_REACHED**
The number of defined Xmirrors on the SLAVE system exceeded the Limit
- **MIRROR_RETRY_OPERATION**
There is an operation in progress on this mirror , please try again your request in a few seconds

Troubleshooting: Please try again the command in a few seconds

- **XMIRROR_SMASTER_INCOMPATIBLE_VERSION**
Xmirror SMaster system version is incompatible for xmirror definition
- **XMIRROR_NUM_OF_XMIRRORS_ON_SMASTER_LIMIT_REACHED**
The number of defined Xmirrors on the SMASTER system exceeded the Limit
- **XMIRROR_SLAVE_INCOMPATIBLE_VERSION**
Xmirror Slave system version is incompatible for xmirror definition
- **MIRRORING_INCOMPATIBLE_TARGET_VERSION**
Mirroring is not supported between the system versions of the specified peers.
- **XMIRROR_NUM_OF_XMIRRORS_ON_MASTER_LIMIT_REACHED**
The number of defined Xmirrors on the MASTER system exceeded the Limit
- **XMIRROR_NAME_ALREADY_EXISTS_ON_MASTER**
Xmirror name already exists on master system, can not create xmirror
- **XMIRROR_NAME_ALREADY_EXISTS_ON_SMASTER**
Xmirror name already exists on smaster system, can not create xmirror
- **XMIRROR_INVALID_MASTER_LOOPBACK_CONFIGURATION**
Cannot define xmirror with Loopback. Loopback target detected on Master system.
- **XMIRROR_INVALID_SMASTER_LOOPBACK_CONFIGURATION**
Cannot define xmirror with Loopback. Loopback target detected on Smaster system.
- **XMIRROR_INVALID_SLAVE_LOOPBACK_CONFIGURATION**
Cannot define xmirror with Loopback. Loopback target detected on Slave system.
- **XMIRROR_MASTER_SMASTER_CLOCK_SKEW_TOO_BIG**
Command failed because the time difference between Master and SMaster systems is too big

Activating an xmirror

Use the **xmirror_activate** command to activate all the mirrors within the xmirror.

xmirror_activate xmirror=XmirrorName

Parameters

Name	Type	Description	Mandatory
xmirror	Object name	xmirror name.	Y

This command activates the xmirror and all of its associated mirrors. When run on the Master, this command activates the Master-Smaster and Master-Slave mirror relations. If one of these relations is already active, nothing is performed for that mirror relation. Running this command on the Smaster or slave is not allowed. Running this command on an already active xmirror is not allowed either.

This command is also used to activate the xmirror after a **change_role** operation. It checks whether the master has more recent data than the slave. If yes, the operation will fail and manual action must be taken.

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed

User Category	Permission
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- **XMIRROR_BAD_NAME**
There is no xmirror with this name
- **XMIRROR_ALREADY_ACTIVE**
xmirror is already active!
- **XMIRROR_ROLE_IS_NOT_MASTER**
Local volume is not the master of the xmirror
- **MIRROR_CONFIGURATION_ERROR**
Mirror local configuration does not match remote configuration
- **MIRROR_CONS_GROUP_MEMBERSHIP_MISMATCH**
Mirrored CG contains different volumes on Master and Slave. This problem occurs whenever the `cg_add_vol` or `cg_remove_vol` commands were previously issued and the Master did not receive an acknowledgment from the Slave until the command timed out, or any other unexpected failure.
- **MIRROR_RETRY_OPERATION**
There is an operation in progress on this mirror , please try again your request in a few seconds

Troubleshooting: Please try again the command in a few seconds

- **MIRROR_SIZE_MISMATCH**
Slave volume and Master Volume sizes are different
- **REMOTE_MAX_COMPRESSED_VOLUMES_REACHED**
Maximum number of compressed volumes already defined in the remote
- **REMOTE_MAX_VOLUMES_REACHED**
Maximum number of volumes already defined on remote machine
- **MAX_VOLUMES_REACHED**
Maximum number of volumes already defined
- **MAX_COMPRESSED_VOLUMES_REACHED**
Maximum number of compressed volumes already defined
- **REMOTE_VOLUME_IS_MASTER**
Volume on remote machine is currently defined as Master
- **XMIRROR_SLAVE_HAS_MORE_RECENT_DATA**
Cannot activate xmirror because slave has more recent data.
- **REMOTE_MIRROR_IS_STANDBY**
remote mirror is marked as standby
- **XMIRROR_SLAVE_INVALID_STANDBY_CONFIGURATION**
Xmirror slave standby mirror configuration is invalid
- **XMIRROR_MASTER_INVALID_CONFIGURATION**
Xmirror master mirror configuration is invalid
- **XMIRROR_NOT_DEFINED**
Xmirror is not defined
- **REMOTE_DOMAIN_MAX_VOLUMES_REACHED**
Maximum number of volumes already defined on remote machine domain
- **XMIRROR_STANDBY_MIRROR_NOT_REGISTERED**
A stand-by mirror not registered or not defined for this xmirror
- **XMIRROR_ASYNC_MIRROR_IS_NOT_CONNECTED**
Cannot activate xmirror after change role if the async mirror is not connected.
- **XMIRROR_MASTER_SLAVE_INCONSISTENT_LRS_CONFIGURATION**
Cannot activate xmirror because master LRS snapshot is more advanced than Slave LRS snapshot. Snapshots configuration is inconsistent.
- **XMIRROR_MASTER_SMASTER_CLOCK_SKEW_TOO_BIG**
Command failed because the time difference between Master and SMaster systems is too big
- **XMIRROR_ROLE_IS_STANDALONE_MASTER**
Command failed because a Xmirror STANDALONE master can not be activated
- **COMPRESSION_VOLUME_IS_OFFLINE**
Operation cannot be completed because the requested compressed volume is offline

Deactivating an xmirror

Use the **xmirror_deactivate xmirror=XmirrorName** command to deactivate all the mirrors in the xmirror relation.

```
xmirror_deactivate xmirror=XmirrorName
```

Parameters

Name	Type	Description	Mandatory
xmirror	Object name	xmirror name.	Y

This command deactivates the xmirror and all of the mirrors in the xmirror relation. When run on the Master, both the Master-Smaster and Master-Slave mirror relations are deactivated. If one of these is already inactive, nothing is performed for that mirror relation. Running this command on the smaster or slave is not allowed.

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_DEACTIVATE_MIRRORING
Are you sure you want to deactivate mirroring?

Return codes

- XMIRROR_BAD_NAME
There is no xmirror with this name
 - XMIRROR_ROLE_IS_NOT_MASTER
Local volume is not the master of the xmirror
 - REMOTE_VOLUME_IS_MASTER
Volume on remote machine is currently defined as Master
 - REMOTE_MIRROR_IS_STANDBY
remote mirror is marked as standby
 - XMIRROR_PARTIAL_DEACTIVATION
Xmirror was partially deactivated
- Troubleshooting:** Deactivate remaining mirrors manually
- XMIRROR_ALREADY_INACTIVE
Xmirror is already inactive.

Deleting an xmirror relation

Use the **xmirror_delete** command to delete an existing 3-way mirror relation.

```
xmirror_delete xmirror=XmirrorName [ force=<yes|no> ]
```

Parameters:

Name	Type	Description	Mandatory	Default
xmirror	Object name	xmirror name.	Y	N/A
force	Boolean	Forces the deletion of the xmirror on the local system.	N	no

Running this command requires that at least one of the Master's mirror relations is active. Running this on the smaster or slave can be done locally by using the force flag, but this will delete the local xmirror only. The force parameter performs the same action as the **mirror_delete** command.

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

- **XMIRROR_BAD_NAME**
There is no xmirror with this name
- **XMIRROR_TOO_MANY_ACTIVE_MIRRORS**
This command cannot be used if more than one mirror is active in the xmirror
- **XMIRROR_ROLE_IS_NOT_MASTER**
Local volume is not the master of the xmirror
- **XMIRROR_DELETE_PARTIAL_FAILURE_MANUAL_CLEANUP_REQUIRED**
Failed to delete xmirror on all systems. Manual cleanup is required

Renaming an xmirror relation

Use the **xmirror_rename** command to rename an xmirror.

```
xmirror_rename xmirror=XmirrorName new_name=Name [ force=<yes|no> ]
```

Parameters

Name	Type	Description	Mandatory	Default
xmirror	Object name	xmirror name.	Y	N/A
new_name	Object name	New xmirror name.	Y	N/A
force	Boolean	If set to yes, forces the rename of the xmirror object on the local system.	N	no

The new name must not exist on any of the systems. Running this command on the smaster or slave can be done using the force flag, in case one of the nodes fails to get renamed.

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

- **XMIRROR_BAD_NAME**
There is no xmirror with this name
- **XMIRROR_ROLE_IS_NOT_MASTER**
Local volume is not the master of the xmirror
- **VOLUME_BAD_PREFIX**
Volume name has a reserved prefix
- **XMIRROR_RENAME_PARTIAL_FAILURE_MANUAL_CLEANUP_REQUIRED**
Failed to rename xmirror. Manual cleanup is required
- **XMIRROR_NAME_ALREADY_EXISTS_ON_MASTER**
Xmirror name already exists on master system, can not create xmirror

Changing roles in an xmirror relation

Use the **xmirror_change_role** command to change a local xmirror peer's role.


```
xmirror_change_role xmirror=XmirrorName new_role=<Master|None|Slave|SMaster>
```

Parameters

Name	Type	Description	Mandatory
xmirror	Object name	xmirror name.	Y
new_role	Enumeration	Role name of the local xmirror.	Y

The following role transitions are allowed:

- Master to Smaster
- Slave to Master
This is effectively "standalone", in the sense that the xmirror relation cannot be redefined on the other systems with the former Slave as the new Master, but rather a new xmirror must be defined.
- Smaster to Master

This command is meant to be used when the couplings are non-operational. It is assumed that the command will be issued on all peers before the couplings become operational again, so that either the original roles will be restored, or the Smaster and Master will have switched roles.

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_PEER_ROLE_TO_SLAVE
Are you sure you want the mirror's local peer to become Slave?
- SOME_DATA_WILL_BE_LOST_ARE_YOU_SURE
Are you sure you want the mirror's local peer to become Slave and loose data that was not replicated?
- ARE_YOU_SURE_YOU_WANT_TO_CHANGE_XMIRROR_ROLE
Are you sure you want to change xmirror role?
- ARE_YOU_SURE_YOU_WANT_TO_CHANGE_XMIRROR_ROLE_WITH_LCS
Are you sure you want to change xmirror role with existing last consistent snapshot?

Return codes

- XMIRROR_BAD_NAME
There is no xmirror with this name
- XMIRROR_SMASTER_CAN_ONLY_BE_CHANGED_TO_MASTER
Xmirror smaster can only be changed to master on this system
- XMIRROR_SLAVE_CAN_ONLY_BE_CHANGED_TO_MASTER
Xmirror slave can only be changed to master on this system
- XMIRROR_MASTER_CAN_ONLY_BE_CHANGED_TO_SLAVE
Xmirror master can only be changed to slave on this system
- MIRROR_RETRY_OPERATION
There is an operation in progress on this mirror , please try again your request in a few seconds
Troubleshooting: Please try again the command in a few seconds
- XMIRROR_MASTER_CAN_ONLY_BE_CHANGED_TO_SMASTER
Xmirror master can only be changed to smaster on this system
- MIRROR_IS_INITIAL
Operation is not permitted during the Initialization phase.
- MIRROR_HAS_NO_SYNCED_SNAPSHOT
Mirror does not have a synchronized Snapshot.
- MASTER_CANNOT_BE_DEMOTED
Master cannot be demoted to Slave role, Peer status mismatch
- MIRROR_IS_ACTIVE
Remote mirroring 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.
- **COMPRESSED_VOLUME_IS_MAPPED**
Mapped compressed volume cannot be a mirror slave
- **COMPRESSION_RECOVERY_IN_PROGRESS**
Operation cannot be completed because another recovery process is in progress

Defining a standby mirror in an xmirror relation

Use the **xmirror_register_standby_mirror** command to define a standby mirror in an existing xmirror relation.

```
xmirror_register_standby_mirror xmirror=XmirrorName slave_target=TargetName
```

Parameters

Name	Type	Description	Mandatory
xmirror	Object name	xmirror name.	Y
slave_target	Object name	Slave target name.	Y

When an xmirror is created with only master-smaster and master-slave mirror relations, this command is used to add the smaster-slave mirror relation to the xmirror. Running this command is allowed on the smaster system only.

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

- **XMIRROR_BAD_NAME**
There is no xmirror with this name
- **TARGET_BAD_NAME**
Target name does not exist
- **XMIRROR_ROLE_IS_NOT_SMASTER**
Local volume is not the smaster of the xmirror
- **XMIRROR_SMASTER_INVALID_CONFIGURATION**
Xmirror substitute-master mirror configuration is invalid
- **XMIRROR_SLAVE_INVALID_STANDBY_CONFIGURATION**
Xmirror slave standby mirror configuration is invalid
- **XMIRROR_STANDBY_MIRROR_ALREADY_DEFINED**
A stand-by mirror is already defined on this system
- **XMIRROR_SLAVE_SMASTER_TARGET_MISMATCH**
Given smaster target does not match defined slave-smaster mirror target
- **XMIRROR_SLAVE_INVALID_CONFIGURATION**
Xmirror slave mirror configuration is invalid
- **XMIRROR_SLAVE_INVALID_MIRROR_STATE**
Invalid mirror state in slave
- **XMIRROR_SMASTER_INVALID_MIRROR_STATE**
Invalid mirror state in substitute master
- **XMIRROR_SMASTER_INVALID_STANDBY_CONFIGURATION**
Invalid standby mirror configuration in smaster system
- **XMIRROR_SMASTER_SLAVE_TARGET_MISMATCH**
Given slave target does not match defined smaster-slave mirror target

- **XMIRROR_INVALID_SMMASTER_LOOPBACK_CONFIGURATION**
Cannot define xmirror with Loopback. Loopback target detected on Smaster system.
- **XMIRROR_INVALID_SLAVE_LOOPBACK_CONFIGURATION**
Cannot define xmirror with Loopback. Loopback target detected on Slave system.

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> ]
```

Parameters

Name	Type	Description	Mandatory	Default
protocol	Enumeration	The available options are: XCLI or All (if no value is specified).	N	All (if no value is specified).

The following default parameters are shown:

- Protocol Name (XCLI)
- Minimal TLS Version (TLS1.0, TLS1.1, TLS1.2)

Example:

```
protocol_config_list
```

Output:

Protocol Name	Minimal TLS Version
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	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

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> ] min_tls_level=<tls1.0|tls1.1|tls1.2>
```

Parameters

Name	Type	Description	Mandatory	Default
protocol	Enumeration	The available options are: XCLI (case insensitive) or All (if no value is specified).	N	All (if no value is specified).

Name	Type	Description	Mandatory	Default
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=XCLI 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

This section contains detailed information on CLI events, including their severity and descriptions.

VOLUME_TRANSFORMATION_QUEUED

Severity	informational
Description	'transform mode' transformation for volume ' <i>volume name</i> ' queued.

VOLUME_TRANSFORMATION_STARTED

Severity	informational
Description	'transform mode' transformation for volume ' <i>volume name</i> ' started.

VOLUME_TRANSFORMATION_SUSPENDED

Severity	informational
Description	'transform mode' transformation for volume ' <i>volume name</i> ' suspended.

VOLUME_TRANSFORMATION_RESUMED

Severity	informational
Description	'transform mode' transformation for volume ' <i>volume name</i> ' resumed.

VOLUME_TRANSFORMATION_FINISHED

Severity	informational
Description	'transform mode' transformation for volume ' <i>volume name</i> ' completed successfully.

VOLUME_TRANSFORMATION_FAILED

Severity	major
Description	'transform mode' transformation for volume ' <i>volume name</i> ' failed due to reason ' <i>failure reason</i> '.

VOLUME_TRANSFORMATION_DELETED

Severity	informational
Description	'transform mode' transformation for volume ' <i>volume name</i> ' deleted/aborted.

COMPRESSION_ENABLED

Severity	informational
Description	Compression was enabled.

COMPRESSION_DISABLED

Severity	informational
Description	Compression was disabled.

COMPRESSION_ENABLING_STARTED

Severity	informational
Description	Compression enabling process has start.

COMPRESSION_DISABLING_STARTED

Severity	informational
Description	Compression disabling has start.

COMPRESSION_DISABLE_IS_TAKING_TOO_LONG

Severity	major
Description	Compression disabling process is taking too long

COMPRESSION_ENABLING_FAILED

Severity	major
Description	Compression enabling process has failed

VOL_COMPRESTIMATE_REQUESTED

Severity	informational
Description	Volume comprestimate requested for volume ' <i>volume name</i> '

VOL_COMPRESTIMATE_COMPLETED

Severity	informational
Description	Volume comprestimate completed for volume ' <i>volume name</i> '

VOL_COMPRESTIMATE_FAILED

Severity	warning
Description	Volume comprestimate failed for volume ' <i>volume name</i> '

VOLUME_CREATE

Severity	informational
Description	Volume was created with name ' <i>volume.name</i> ' and size <i>volume.sizeGB</i> in Storage Pool with name ' <i>volume.pool_name</i> '.

VOLUME_CREATE_MANY

Severity	informational
Description	<i>number</i> Volumes was created with names: ' <i>names</i> ' in Storage Pool with name ' <i>pool.name</i> '.

VOLUME_CREATE_FAILED_TOO_MANY

Severity	warning
Description	Volume with name ' <i>name</i> ' could not be created. You are attempting to add more volumes than the system permits.
Troubleshooting	Delete volumes to allow new ones to be created.

VOLUME_CREATE_FAILED_BAD_SIZE

Severity	warning
Description	Volume with name ' <i>name</i> ' could not be created with size of <i>requested_sizeGB</i> . Volume size is not a multiple of the volume size quanta (Partitions).
Troubleshooting	Set volume size that is an integer multiple of 16K (number of slices) partitions.

VOLUME_RENAME

Severity	informational
Description	Volume with name ' <i>old_name</i> ' and was renamed ' <i>volume.name</i> '.

VOLUME_RESIZE

Severity	informational
Description	Volume with name ' <i>volume.name</i> ' was resized from <i>old_size</i> GB to <i>volume.size</i> GB.

VOL_SET_EXTERNAL_ID

Severity	informational
Description	Volume with name ' <i>volume.name</i> ' changed the external identifier to ' <i>volume.identifier</i> '.

VOL_CLEAR_EXTERNAL_ID

Severity	informational
Description	Volume with name ' <i>volume.name</i> ' cleared the external identifier.

SECONDARY_VOLUME_RESIZE

Severity	informational
Description	Secondary volume with name ' <i>volume.name</i> ' was resized by primary machine from <i>old_size</i> GB to <i>volume.size</i> GB.

VOLUME_DELETE

Severity	informational
Description	Volume with name ' <i>volume.name</i> ' was deleted and its data is no longer accessible.

VOLUME_FORMAT

Severity	informational
Description	Volume with name ' <i>volume.name</i> ' was formatted.

VOLUME_COPY

Severity	informational
Description	Volume with name ' <i>source.name</i> ' was copied to volume with name ' <i>target.name</i> '.

VOLUME_COPY_DIFF

Severity	informational
Description	Volume with name ' <i>source.name</i> ' was diff-copied from base ' <i>base.name</i> ' to volume with name ' <i>target.name</i> '.

VOLUME_LOCK

Severity	informational
Description	Volume with name ' <i>volume.name</i> ' was locked and set to 'read-only'.

VOLUME_UNLOCK

Severity	informational
Description	Volume with name ' <i>volume.name</i> ' was unlocked and set to 'writable'.

VOLUME_MOVE

Severity	informational
Description	Volume with name ' <i>volume.name</i> ' has been moved from Storage Pool ' <i>orig_pool.name</i> ' to Pool ' <i>pool.name</i> '.

VOLUME_UNFORMAT

Severity	informational
Description	Volume with name ' <i>volume.name</i> ' was unformatted.

VOLUME_SET_FLASH_BYPASS

Severity	informational
Description	Flash Cache Bypass was set to be ' <i>Bypass</i> ' for Volume with name ' <i>volume.name</i> '.

VOLUME_SET_SSD_CACHING

Severity	informational
Description	SSD Caching was set to be ' <i>state</i> ' for Volume with name ' <i>volume.name</i> '.

VOLUME_SET_ALL_SSD_CACHING

Severity	informational
Description	SSD Caching was set to be ' <i>state</i> ' for all currently defined Volumes.

VOLUME_SET_DEFAULT_SSD_CACHING

Severity	informational
Description	Default SSD Caching for volumes was set to be ' <i>state</i> '.

OLVM_CREATE

Severity	informational
Description	Volume with name ' <i>volume.name</i> ' was defined as an IBM Hyper-Scale Mobility.

OLVM_OWNER_CREATE

Severity	informational
Description	Volume with name ' <i>volume.name</i> ' was defined as an OWNER IBM Hyper-Scale Mobility.

OLVM_OWNER_ACTIVATED

Severity	informational
Description	IBM Hyper-Scale Mobility Owner Volume with name ' <i>volume.name</i> ' was activated.

OLVM_ACTIVATE

Severity	informational
Description	IBM Hyper-Scale Mobility Volume with name ' <i>volume.name</i> ' was activated.

OLVM_DEACTIVATE

Severity	informational
Description	IBM Hyper-Scale Mobility Volume with name ' <i>volume.name</i> ' was deactivated.

OLVM_PROXY_INITIATED

Severity	informational
Description	IBM Hyper-Scale Mobility Volume with name ' <i>volume.name</i> ' move to Proxy state started.

OLVM_PROXY

Severity	informational
Description	IBM Hyper-Scale Mobility Volume with name ' <i>volume.name</i> ' entered Proxy state. Volume data on the system is freed.

OLVM_DELETE

Severity	informational
Description	IBM Hyper-Scale Mobility Volume process with name ' <i>name</i> ' was deleted.
Troubleshooting	Delete volumes to allow new ones to be created.

OLVM_ABORT

Severity	informational

Description	IBM Hyper-Scale Mobility Volume process with name ' <i>volume.name</i> ' was aborted.

OLVM_OWNER_DELETE

Severity	informational
Description	IBM Hyper-Scale Mobility Owner Volume process with name ' <i>volume.name</i> ' was deleted.

OLVM_OWNER_ABORT

Severity	informational
Description	IBM Hyper-Scale Mobility Owner Volume process with name ' <i>volume.name</i> ' was aborted.

OLVM_SYNC_STARTED

Severity	informational
Description	IBM Hyper-Scale Mobility Synchronization of volume ' <i>name</i> ' has started.

OLVM_SYNC_ENDED

Severity	informational
Description	IBM Hyper-Scale Mobility Synchronization of volume ' <i>name</i> ' has ended.

OLVM_BLACKLIST_EDITED

Severity	informational
Description	Volume serial ' <i>serial</i> ' was cleared from the blacklist.

OLVM_BLACKLIST_CLEARED

Severity	major
Description	All volume serials were cleared from the blacklist.

OLVM_BLACKLIST_FULL

Severity	major
Description	Cannot create new IBM Hyper-Scale Mobility Relations. Too many volume serials are blacklisted.

OLVM_LIMITS_CHANGED

Severity	major
Description	IBM Hyper-Scale Mobility limits were changed. Maximum user volumes now at ' <i>max_user_volumes</i> ', blacklist limit at ' <i>blacklist_limit</i> '.

COMPRESSED_VOLUME_LOCK

Severity	Major
Description	Volume ' <i>volume_name</i> ' was locked. Lock type: <i>lock type</i>

COMPRESSED_VOLUME_OFFLINE

Severity	major
Description	Volume ' <i>volume_name</i> ' is offline due to metadata corruption. Volume was locked to all IOs.

COMPRESSED_VOLUME_ONLINE

Severity	informational
Description	Volume ' <i>volume_name</i> ' is online due to vol_unlock command after recovery successful.

COMPRESSED_VOLUME_FORCED_ONLINE

Severity	informational
Description	Volume ' <i>volume_name</i> ' is online due to force online command.

COMPRESSED_VOLUME_VERSION_MISMATCH

Severity	major
Description	The compression format version of volume ' <i>volume_name</i> ' is not supported by this system. Please contact technical support.

COMPRESSED_VOLUME_RECOVERY_STARTED

Severity	major
Description	Recovery process has started on volume ' <i>volume_name</i> ', recovery mode: ' <i>recovery_mode</i> '

COMPRESSED_VOLUME_RECOVERY_FINISHED

Severity	major
Description	Finished recovery process on volume ' <i>volume_name</i> '

COMPRESSED_VOLUME_RECOVERY_ABORTED

Severity	major
Description	Recovery process on volume ' <i>volume_name</i> ' was aborted

COMPRESSED_VOLUME_MARGINAL_COMPRESSION

Severity	variable
Description	Volume ' <i>volume_name</i> ' has a marginal compression and might get locked, consider resizing the volume.

COMPRESSED_VOLUME_ACCEPTABLE_COMPRESSION

Severity	informational
Description	Volume ' <i>volume_name</i> ' is back to acceptable compression.

COMPRESSED_VOLUME_TRIM

Severity	informational
Description	Compressed volume ' <i>name</i> ' trim was triggered. last trim size: <i>last_trim_size</i> , current: <i>current_trim_size</i>

DATA_REBUILD_STARTED

Severity	informational
Description	Rebuild process started because system data is not protected. <i>data_percent</i> % of the data must be rebuilt.

DATA_REBUILD_COMPLETED

Severity	informational
Description	Rebuild process completed. System data is now protected.

DATA_REBUILD_COULD_NOT_BE_COMPLETED

Severity	major
Description	Rebuild process could not be completed due to insufficient unused disk space. System data is not protected.
Troubleshooting	Replace failed drives, delete unused pools or decrease pool size where possible.

DATA_REDIST_STARTED

Severity	informational
Description	Starting data transfer to new disks.

OPTIMIZING_DATA_REDIST_STARTED

Severity	informational
Description	Starting optimizing data transfer to new disks.

DATA_REDIST_COMPLETED

Severity	informational
Description	Completed data transfer to new disks.

DATA_REBUILD_COMPLETED_REDIST_STARTED

Severity	informational
Description	Rebuild process completed. System data is now protected. Starting data transfer to new disks.

STORAGE_POOL_EXHAUSTED

Severity	major
Description	Pool ' <i>pool</i> ' is full. All volumes are locked.
Troubleshooting	Enlarge Storage Pool or move or delete volumes or Clones with Clone Deletion Priority 0.

STORAGE_POOL_CONTAINING_COMPRESSED_VOLUMES_EXHAUSTED

Severity	major
Description	Pool ' <i>pool</i> ' is full. The compressed volumes of the pool are locked for both read and write IOs.
Troubleshooting	Enlarge Storage Pool or move or delete volumes or Clones with Clone Deletion Priority 0.

STORAGE_POOL_UNLOCKED

Severity	major
Description	Pool ' <i>pool</i> ' has empty space. All volumes are unlocked.
Troubleshooting	N/A

STORAGE_POOL_VOLUME_USAGE_INCREASED

Severity	variable
Description	Usage by volumes of Storage Pool with name ' <i>pool.name</i> ' has reached <i>current%</i> .
Troubleshooting	N/A

STORAGE_POOL_VOLUME_USAGE_DECREASED

Severity	informational
Description	Usage by volumes of Storage Pool with name ' <i>pool.name</i> ' has decreased to <i>current%</i> .
Troubleshooting	N/A

STORAGE_POOL_VOLUME_USAGE_TOO_HIGH

Severity	major
Description	Usage by volumes of Storage Pool with name ' <i>pool.name</i> ' has reached <i>current%</i> of the total pool size.
Troubleshooting	Increase pool size or decrease snapshot size.

STORAGE_POOL_VOLUME_USAGE_BACK_TO_NORMAL

Severity	informational
Description	Usage by volumes of Storage Pool with name ' <i>pool.name</i> ' is back to normal with <i>current%</i> of the total pool size.
Troubleshooting	N/A

STORAGE_POOL_SNAPSHOT_USAGE_INCREASED

Severity	variable
Description	Usage by snapshots of Storage Pool with name ' <i>pool.name</i> ' has reached <i>current%</i> .
Troubleshooting	N/A

STORAGE_POOL_SNAPSHOT_USAGE_DECREASED

Severity	informational
Description	Usage by snapshots of Storage Pool with name ' <i>pool.name</i> ' has decreased to <i>current%</i> .
Troubleshooting	N/A

HOST_CONNECTED

Severity	informational
Description	Host ' <i>host</i> ' has connected to the system.
Troubleshooting	N/A

HOST_DISCONNECTED

Severity	warning
Description	Host ' <i>host</i> ' has disconnected from the system.
Troubleshooting	N/A

HOST_MULTIPATH_OK

Severity	informational
Description	Host ' <i>host</i> ' has redundant connections to the system. <i>#paths=npaths</i>
Troubleshooting	N/A

HOST_NO_MULTIPATH_ONLY_ONE_PORT

Severity	warning
Description	Host ' <i>host</i> ' is connected to the system through only one of its ports. <i>#paths=npaths</i>
Troubleshooting	N/A

HOST_NO_MULTIPATH_ONLY_ONE_MODULE

Severity	informational
Description	Host ' <i>host</i> ' is connected to the system through only one interface module. <i>#paths=npaths</i>
Troubleshooting	N/A

SYSTEM_SPARES_ARE_LOW

Severity	major
Description	System capacity spares are <i>modules</i> modules and <i>disks</i> disks.
Troubleshooting	N/A

SYSTEM_NO_SPARES

Severity	critical
Description	System has no spare disks
Troubleshooting	N/A

CMP_VERY_LONG_LATENCY_TO_CACHE

Severity	variable
Description	Compression <i>#Node</i> has long latencies to cache node <i>cache</i> for <i>interval</i> secs.
Troubleshooting	Please contact support.

CMP_LONG_LATENCY_TO_CACHE

Severity	variable
Description	Compression <i>#Node</i> had long latencies to cache node <i>cache</i> for <i>interval</i> msecs.
Troubleshooting	Please contact support.

COMPRESSED_HANDLE_IS_OFFLINE

Severity	major
Description	RACE Handle <i>Handle</i> of volume <i>Volume</i> has corrupted metadata and moved to offline mode

COMPRESSED_HANDLE_SDL_THRESHOLD_IS_REACHED

Severity	major
Description	RACE Handle <i>Handle</i> of volume <i>Volume</i> has crossed the SDL threshold with <i>Errors</i> errors and goes offline.

CMP_RACE_MD_IO_SHOULD_HAVE_TIMED_OUT

Severity	warning
Description	RACE MD IO to handle <i>0xHandle</i> of volume <i>0xVolume</i> should have been timed out if it was not MD

COMPRESSION_NODE_FAILED_CREATING_NETCOMP_CFG_FILE

Severity	Major
Description	Compression node <i>Node ID</i> failed creating NetComp.cfg file

CMP_RACE_ASSERT

Severity	Critical
Description	RACE library on compression node <i>Node ID</i> asserted.

CMP_PARTIAL_WRITE

Severity	warning
Description	Compression node <i>#Node</i> reports partial write to volume ' <i>volume</i> ' on lba <i>lba</i> , primary failed= <i>primary_failed</i> , secondary failed= <i>secondary_failed</i> , remote failed= <i>remote_failed</i> .

POOL_CREATE

Severity	informational
Description	Storage Pool of size <i>pool.sizeGB</i> was created with name ' <i>pool.name</i> '.

POOL_CREATE_THIN

Severity	informational
Description	Storage Pool of soft size <i>pool.soft_sizeGB</i> and hard_size <i>pool.hard_sizeGB</i> was created with name ' <i>pool.name</i> '.

POOL_CREATE_FAILED_TOO_MANY

Severity	warning
Description	Storage Pool with name ' <i>name</i> ' could not be created. You are attempting to add more Storage Pools than the system permits.
Troubleshooting	Delete Storage Pools to allow new ones to be created.

POOL_RENAME

Severity	informational
Description	Storage Pool with name ' <i>old_name</i> ' was renamed ' <i>pool.name</i> '.

POOL_RESIZE

Severity	informational
Description	Storage Pool with name ' <i>pool.name</i> ' was resized from size <i>old_sizeGB</i> to <i>pool.sizeGB</i> .

POOL_RESIZE_THIN

Severity	informational
Description	Storage Pool with name ' <i>pool.name</i> ' was resized from soft size <i>old_soft_sizeGB</i> and hard size <i>old_hard_sizeGB</i> to soft size <i>pool.soft_sizeGB</i> and hard size <i>pool.hard_sizeGB</i> .

POOL_RESIZE_SNAPSHOTS

Severity	informational
Description	Snapshot size of Storage Pool with name ' <i>pool.name</i> ' was resized from size <i>old_sizeGB</i> to <i>pool.snapshot_sizeGB</i> .

POOL_CHANGE_LOCK_BEHAVIOR

Severity	informational
Description	Lock Behavior of Storage Pool with name ' <i>pool.name</i> ' is now ' <i>state</i> '.

POOL_CHANGE_PERF_CLASS

Severity	informational
Description	Performance Class of Storage Pool with name ' <i>pool.name</i> ' is now ' <i>pool.perf_class</i> '.

POOL_CONFIG_SNAPSHOTS

Severity	informational
Description	Management policy of Mirroring snapshots of Storage Pool with name ' <i>pool.name</i> ' has changed'.

POOL_DELETE

Severity	informational
Description	Storage Pool with name ' <i>pool.name</i> ' was deleted.

COMMAND_SERVICE_FAILED_TOO_MANY_TIMES

Severity	major
Description	Command service <i>service name</i> was restarted <i>times</i> times within <i>seconds</i> seconds on module <i>module</i>
Troubleshooting	Please contact support.

COMMAND_SERVICE_EXECUTABLE_INACCESSIBLE

Severity	major
Description	Command service <i>service name</i> 's executable was not found on module <i>module</i>
Troubleshooting	Please contact support.

FC_PORT_RESTART

Severity	variable
Description	FC port service <i>port</i> was restarted due to <i>code codestr</i>
Troubleshooting	Please contact support.

ISCSI_PORT_RESTART

Severity	variable
Description	ISCSI port service <i>port</i> was restarted due to <i>code codestr</i>
Troubleshooting	Please contact support.

FC_PORT_HAS_FAILED

Severity	major
Description	FC port service <i>port</i> has failed due to <i>code codestr</i> (attempt number <i>Number of retries</i>)
Troubleshooting	Please contact support.

ISCSI_PORT_HAS_FAILED

Severity	major
Description	ISCSI port service <i>port</i> has failed due to <i>code codestr</i> (attempt number <i>Number of retries</i>)
Troubleshooting	Please contact support.

REQUIREMENT_IS_MISSING

Severity	critical
Description	Requirement <i>requirement name</i> is missing

CRITICAL_THREAD_DID_NOT_HEARTBEAT

Severity	critical
Description	Thread named ' <i>thread name</i> ' in process ' <i>process_name</i> ' for node <i>node id</i> on module <i>module</i> did not heartbeat
Troubleshooting	Please contact support.

THREAD_DID_NOT_HEARTBEAT

Severity	warning
Description	Thread named ' <i>thread name</i> ' in process ' <i>process_name</i> ' for node <i>node id</i> on module <i>module</i> did not heartbeat
Troubleshooting	Please contact support.

NODE_DID_NOT_HEARTBEAT

Severity	warning
Description	Node named ' <i>process_name</i> ' with id <i>node id</i> on module <i>module</i> did not heartbeat
Troubleshooting	Please contact support.

MODULE_STARTED_DOWNLOADING_VERSION

Severity	informational
Description	Module <i>Module ID</i> started downloading current version of the system

MODULE_FINISHED_DOWNLOADING_VERSION

Severity	variable
Description	Module <i>Module ID</i> finished downloading current version of the system. Downloaded total of <i>Number of files</i> files. Status: <i>Status</i>

MODULE_DOWNLOAD_TIMEOUT

Severity	warning
Description	Timeout expired trying to download current version of the system to module <i>Module ID</i> using interface <i>Interface</i> .
Troubleshooting	Please contact support.

MODULE_DOWNLOAD_VERSION_TIMEOUT

Severity	warning
Description	Timeout expired trying to download current version of the system to module <i>Module ID</i> .
Troubleshooting	Please contact support.

ERROR_SETTING_UP_INTERFACE

Severity	major
Description	Error attempting setup Ethernet interface <i>Interface name</i> on module <i>Module ID</i>
Troubleshooting	Possible physical problem with Ethernet card. Contact support

MODULE_CPU_HAS_LESS_CORES_THAN_EXPECTED

Severity	major
Description	CPU of <i>Component ID</i> has less cores than expected: got <i>actual cores</i> , expected <i>req cores</i> .
Troubleshooting	Please contact your Administrator.

MODULE_CPU_HAS_MORE_CORES_THAN_EXPECTED

Severity	informational
Description	CPU of <i>Component ID</i> has more cores than expected: got <i>actual cores</i> cores, expected only <i>req cores</i> .
Troubleshooting	Please contact your Administrator.

MISMATCH_IN_INTERFACE_SPEED

Severity	major
Description	Interface speed on <i>Component ID</i> is <i>actual speedG</i> , the expected speed is <i>req speedG</i> .
Troubleshooting	Please contact your Administrator.

NEW_TIME_CHANGE_IS_INVALID

Severity	warning
Description	Setting time to <i>Seconds</i> seconds and <i>USecs</i> Useconds on module <i>Module</i> is invalid and was denied.
Troubleshooting	Please contact your Administrator.

NTP_SERVER_TIME_DIFFERENCE_TOO_BIG

Severity	warning
Description	NTP server <i>NTP Server</i> sent a transaction with time difference of <i>Delta</i> seconds which exceeds the maximal difference of <i>Max Allowed</i> seconds. Transaction will be ignored, please check NTP server's and system's times.
Troubleshooting	Please contact your Administrator.

IPSEC_TUNNEL_OPENED

Severity	informational
Description	The IPSec tunnel named ' <i>name</i> ' between module <i>Module</i> and <i>Right IP</i> was opened

IPSEC_TUNNEL_CLOSED

Severity	informational
Description	The IPSec tunnel named ' <i>name</i> ' between module <i>Module</i> and <i>Right IP</i> was closed

MODULE_ROOT_FILESYSTEM_IS_MOUNTED_READ_WRITE

Severity	minor
Description	The root file system of <i>Component ID</i> is mounted as read-write, it's impossible to check it for corruption.
Troubleshooting	Please contact support.

IP_ACCESS_CANNOT_RESOLVE_ADDRESS

Severity	informational
Description	Cannot resolve address ' <i>address</i> ' added to the IP access group <i>IP access group name</i> .

IP_ACCESS_FAILED_SETTING_RULES

Severity	informational
Description	Failed setting IP access rules.

CACHE_HAS_LESS_MEMORY

Severity	warning
Description	Data module has less memory than expected. node= <i>node</i> - <i>gb_missing</i> GB missing.
Troubleshooting	some of the DIMMs might have failed

LOAD_MONITOR_NATIVE_PROCESS_MEMORY_USAGE_CHANGED

Severity	variable
Description	Process ' <i>Process Name</i> ' with PID <i>PID</i> on module <i>Module</i> . Threshold: <i>Threshold Type</i> . Memory usage: <i>Memory usageKb</i> .

LOAD_MONITOR_NATIVE_PROCESS_FD_USAGE_CHANGED

Severity	variable
Description	Process ' <i>Process Name</i> ' with PID <i>PID</i> on module <i>Module</i> . Threshold: <i>Threshold Type</i> . FD usage: <i>FD usage</i>

LOAD_MONITOR_NATIVE_PROCESS_CPU_USAGE_CHANGED

Severity	variable
Description	Process ' <i>Process Name</i> ' with PID <i>PID</i> on module <i>Module</i> . Threshold: <i>Threshold Type</i> . CPU usage: <i>CPU usage</i>

LOAD_MONITOR_NUMBER_OF_ALIEN_PROCESSES_EXCEEDS_THRESHOLD

Severity	warning
Description	Number of alien processes on module <i>Module</i> exceeds limit. Threshold <i>Threshold</i>

LOAD_MONITOR_ALIEN_PROCESSES_MEMORY_USAGE_CHANGED

Severity	variable
Description	Alien Processes on module <i>Module</i> . Threshold: <i>Threshold Type</i> . Memory usage: <i>Memory usageKb</i> .

LOAD_MONITOR_ALIEN_PROCESSES_FD_USAGE_CHANGED

Severity	variable
Description	Alien Processes on module <i>Module</i> . Threshold: <i>Threshold Type</i> . FD usage: <i>FD usage</i>

LOAD_MONITOR_ALIEN_PROCESSES_CPU_USAGE_CHANGED

Severity	variable
Description	Alien Processes on module <i>Module</i> . Threshold: <i>Threshold Type</i> . CPU usage: <i>CPU usage</i>

LOAD_MONITOR_PROCESS_KILLED

Severity	warning
Description	Process ' <i>Process Name</i> ' with PID <i>PID</i> on module <i>Module</i> was killed. Reason: <i>Killing reason</i> . Usage: <i>Usage</i>

LOAD_MONITOR_PROCESS_KILL_FAILED

Severity	minor
Description	Process ' <i>Process Name</i> ' with PID <i>PID</i> on module <i>Module</i> was not killed. Reason: <i>Killing reason</i> . Usage: <i>Usage</i>

LOAD_MONITOR_TOTAL_COMMITTED_MEMORY_DELTA_CHANGED

Severity	variable
Description	Total committed memory delta on module <i>Module</i> . Threshold: <i>Threshold Type</i> . Total committed memory delta: <i>Total committed memory delta</i>

LOAD_MONITOR_TOTAL_SLAB_USAGE_CHANGED

Severity	variable
Description	Total slab usage on module <i>Module</i> . Threshold: <i>Threshold Type</i> . Total slab usage: <i>Total slab usage</i>

LOAD_MONITOR_NATIVE_SLAB_USAGE_CHANGED

Severity	variable
Description	Slab ' <i>Slab Name</i> ' on module <i>Module</i> . Threshold: <i>Threshold Type</i> . Usage: <i>Slab usageKb</i> .

LOAD_MONITOR_MOUNTING_POINT_USED_SPACE_CHANGED

Severity	variable
Description	Used space on mounting point ' <i>mounting_point_path</i> ' on module <i>Module</i> has changed. Threshold: <i>Threshold Type</i> . Used space: <i>Used space%</i> .

LOAD_MONITOR_ALIEN_SLAB

Severity	minor
Description	Alien slab ' <i>Slab Name</i> ' on module <i>Module</i> . Usage: <i>Slab usageKb</i> .

LOAD_MONITOR_NATIVE_IRQ_USAGE_CHANGED

Severity	variable
Description	IRQ from ' <i>Device Name</i> ' type <i>Type</i> number <i>IRQ</i> on module <i>Module</i> . Threshold: <i>Threshold Type</i> . Usage: <i>Usage</i> .

LOAD_MONITOR_ALIEN_IRQ

Severity	minor
Description	Alien IRQ from ' <i>Device Name</i> ' type <i>Type</i> on module <i>Module</i> . Usage: <i>Usage</i> .

ROOT_RW_REMOUNT_TIMEOUT

Severity	minor
Description	Root R/W remount with id ' <i>ID</i> ' has timed out.

SAS_CONTROLLER_DIED

Severity	warning
Description	Severe SAS controller error occurred. Controller was removed from PCI-E bus.
Troubleshooting	Please contact support.

SAS_CONTROLLER_IMPLICIT_RESET_SUCCESSFUL

Severity	warning
Description	SAS driver sent an implicit reset to SAS controller, controller was successfully reset.
Troubleshooting	Please contact support.

SAS_CONTROLLER_IMPLICIT_RESET_FAILED

Severity	warning
Description	SAS driver sent an implicit reset to SAS controller, but it failed.
Troubleshooting	Please contact support.

SAS_CONTROLLER_BIGHAMMER_SUCCESSFUL

Severity	warning
Description	SAS controller was successfully reset with 'BIG HAMMER' reset.
Troubleshooting	Please contact support.

SAS_CONTROLLER_BIGHAMMER_FAILED

Severity	warning
Description	SAS controller 'BIG HAMMER' reset was attempted, but it failed.
Troubleshooting	Please contact support.

MISSED_KERNEL_EVENTS

Severity	warning
Description	Some kernel events were missed.
Troubleshooting	Please contact support.

BOIDEM_DISK_DEFERRED_ERROR

Severity	warning
Description	Deferred error on <i>Disk ID</i> , start LBA= <i>Start LBA</i> , last LBA= <i>Last LBA</i> , latency= <i>latency</i> msec, key= <i>key</i>
Troubleshooting	N/A

BOIDEM_DISK_REVIVED

Severity	warning
Description	Boidem disk <i>Disk ID</i> revived.

BOIDEM_DISK_KILLED

Severity	warning
Description	Boidem disk <i>Disk ID</i> killed.

BOIDEM_DISK_BLACKLIST_MASK_CHANGED

Severity	informational
Description	Boidem disk blacklist changed on module <i>module</i> from <i>old_blacklist</i> to <i>new_blacklist</i> .

BOIDEM_NOT_ENOUGH_DISKS_AFTER_GRACE

Severity	warning
Description	Grace period expired, but there are still not enough disks for boidem RAID on module <i>module</i> .

BOIDEM_HAS_ENOUGH_DISKS

Severity	warning
Description	There are now enough disks for boidem RAID on module <i>module</i> .

BOIDEM_MISSING_MOUNT_POINT

Severity	warning
Description	Boidem is missing a mount point at <i>Missing mount point</i> on module <i>module</i> .

BOIDEM_FS_IS_RO

Severity	warning
Description	Boidem mount point <i>Read-only mount point</i> is in a read-only state on module <i>module</i> .

BOIDEM_DISK_UNRESPONSIVE

Severity	warning
Description	Disk <i>Disk ID</i> is unresponsive for <i>time</i> ms

BOIDEM_DISK_RESPONSIVE

Severity	warning
Description	Disk <i>Disk ID</i> is now responsive. Was unresponsive for <i>unresponsive_time</i> ms

BOIDEM_DISK_ERROR_SENSE_INFORMATION

Severity	minor
Description	Disk <i>Disk ID</i> had sense information indicating an error: <i>Sense Key Number/Sense Code Number 1/Sense Code Number 2 (FRU=FRU Code) Sense Key - Sense Code</i> .
Troubleshooting	Comes together with other disk errors

BOIDEM_DISK_MEDIUM_ERROR

Severity	warning
Description	Media errors on <i>Disk ID</i> , start LBA= <i>Start LBA</i> , last LBA= <i>Last LBA</i> , latency= <i>latency</i> ms.

BOIDEM_DISK_ABNORMAL_ERROR

Severity	warning
Description	Unit attentions or aborts in the last 30 minutes on <i>Disk ID</i> , start lba= <i>start_lba</i> , last lba= <i>last_lba</i> , command= <i>command</i> , latency= <i>latency</i> ms.

BOIDEM_DISK_LONG_LATENCY

Severity	variable
Description	Disk <i>Disk ID</i> has been exhibiting long I/O latency in the last 30 minutes, start LBA= <i>Start LBA</i> , last LBA= <i>Last LBA</i> , command= <i>command</i> , latency= <i>latency</i> ms.

BOIDEM_DISK_KEEPAIVE_OK

Severity	minor
Description	Disk <i>Disk ID</i> is responding to keepalives of type <i>Type</i> after <i>Time from last success msecms</i>

BOIDEM_DISK_KEEPAIVE_FAILED

Severity	warning
Description	Disk <i>Disk ID</i> is not responding to keepalives of type <i>Type</i> for <i>Time from last success msecms</i>
Troubleshooting	N/A

DISK_KEEPAIVE_FAILED

Severity	major
Description	Disk <i>Disk ID</i> is not responding to keepalives of type <i>Type</i> for <i>Time from last success msecms</i> , disk cache dirty level is <i>Disk cache dirty level%</i>
Troubleshooting	N/A

SSD_KEEPAIVE_FAILED

Severity	major
Description	SSD <i>SSD ID</i> is not responding to keepalives of type <i>Type</i> for <i>Time from last success msecms</i>
Troubleshooting	N/A

DISK_KEEPAIVE_OK

Severity	major
Description	Disk <i>Disk ID</i> is responding to keepalives of type <i>Type</i> after <i>Time from last success msecms</i>
Troubleshooting	N/A

SSD_KEEPAIVE_OK

Severity	major
Description	SSD <i>SSD ID</i> is responding to keepalives of type <i>Type</i> after <i>Time from last success msecms</i>
Troubleshooting	N/A

BUS_RESET_WAS_SENT

Severity	warning
Description	A bus reset was sent on module <i>module</i> .
Troubleshooting	Please contact support.

BUS_RESET_FAILED

Severity	major
Description	Bus reset on module <i>module</i> has failed. Reset duration <i>reset duration</i> , IOs pending <i>IOs Pending</i> .
Troubleshooting	Please contact support.

BUS_RESET_SUCCEEDED

Severity	informational
Description	Bus reset on module <i>module</i> succeeded. Reset duration <i>reset duration</i> , IOs pending <i>IOs Pending</i> .
Troubleshooting	Please contact support.

HOST_RESET_WAS_SENT

Severity	warning
Description	A host reset was sent on module <i>module</i> .
Troubleshooting	Please contact support.

HOST_RESET_FAILED

Severity	major
Description	Host reset on module <i>module</i> has failed. Reset duration <i>reset duration</i> , IOs pending <i>IOs Pending</i> .
Troubleshooting	Please contact support.

HOST_RESET_SUCCEEDED

Severity	informational
Description	Host reset on module <i>module</i> succeeded. Reset duration <i>reset duration</i> , IOs pending <i>IOs Pending</i> .
Troubleshooting	Please contact support.

DISK_RESET_WAS_SENT

Severity	warning
Description	A disk reset was sent to <i>Component ID</i> .
Troubleshooting	Please contact support.

DISK_RESET_FAILED

Severity	major
Description	Reset to disk <i>Component ID</i> has failed. Reset duration <i>reset duration</i> , IOs pending <i>IOs Pending</i> .
Troubleshooting	Please contact support.

DISK_RESET_SUCCEEDED

Severity	informational
Description	Reset to disk <i>Component ID</i> succeeded. Reset duration <i>reset duration</i> , IOs pending <i>IOs Pending</i> .
Troubleshooting	Please contact support.

SAS_CONTROLLER_RESET_WAS_SENT

Severity	warning
Description	A SAS controller reset was sent on <i>Component ID</i> , IOs pending <i>IOs Pending</i> .
Troubleshooting	Please contact support.

SAS_CONTROLLER_RESET_FAILED

Severity	major
Description	Reset to the SAS controller on <i>Component ID</i> has failed. Reset duration <i>reset duration</i> milliseconds, IOs pending <i>IOs Pending</i> .
Troubleshooting	Please contact support.

SAS_CONTROLLER_RESET_SUCCEEDED

Severity	informational
Description	Reset to disk <i>Component ID</i> succeeded. Reset duration <i>reset duration</i> milliseconds, IOs pending <i>IOs Pending</i> .
Troubleshooting	Please contact support.

DISK_RESET_FAILURE

Severity	major
Description	Reset to disk <i>Component ID</i> was executed and failed. Reset duration <i>reset duration</i> usecs, IOs pending <i>IOs Pending</i> .
Troubleshooting	Please contact support.

SSD_RESET_FAILURE

Severity	major
Description	Reset to disk <i>Component ID</i> was executed and failed. Reset duration <i>reset duration</i> usecs, IOs pending <i>IOs Pending</i> .
Troubleshooting	Please contact support.

DISK_RESET_DONE

Severity	warning
Description	Reset to disk <i>Component ID</i> was executed and succeeded. Reset duration <i>reset duration</i> usecs, IOs pending <i>IOs Pending</i> .
Troubleshooting	Please contact support.

SSD_RESET_DONE

Severity	warning
Description	Reset to disk <i>Component ID</i> was executed and succeeded. Reset duration <i>reset duration</i> usecs, IOs pending <i>IOs Pending</i> .
Troubleshooting	Please contact support.

DISK_POWER_DOWN

Severity	major
Description	Disk <i>Component ID</i> was powered-down due to error recovery failures.
Troubleshooting	Please contact support.

SSD_OFFLINE

Severity	major
Description	SSD <i>Component ID</i> was marked as offline due to error recovery failures.
Troubleshooting	Please contact support.

DISK_ERROR_SENSE_INFORMATION

Severity	minor
Description	Disk <i>Disk ID</i> had sense information indicating an error: <i>Sense Key Number/Sense Code Number 1/Sense Code Number 2</i> (FRU=FRU Code) <i>Sense Key - Sense Code</i> .
Troubleshooting	Comes together with other disk errors

SSD_ERROR_SENSE_INFORMATION

Severity	minor
Description	SSD <i>SSD ID</i> had sense information indicating an error: <i>Sense Key Number/Sense Code Number 1/Sense Code Number 2</i> (FRU=FRU Code) <i>Sense Key - Sense Code</i> .
Troubleshooting	Comes together with other error events

DISK_REQUEST_ERROR_INFORMATION

Severity	minor
Description	Disk <i>Disk ID</i> had error: <i>Error Name</i> , latency= <i>latency</i> ms.
Troubleshooting	Comes together with other disk errors

SSD_REQUEST_ERROR_INFORMATION

Severity	minor
Description	SSD <i>SSD ID</i> had error: <i>Error Name</i> , latency= <i>latency</i> ms.
Troubleshooting	Comes together with other error events

SSD_RECOVERED_ERROR

Severity	minor
Description	SSD <i>SSD ID</i> autonomously recovered from an error successfully, start lba= <i>first_lba</i> , last lba= <i>last_lba</i> , scsi_opcode= <i>scsi_opcode</i> , latency= <i>latency</i> usec.
Troubleshooting	N/A

DISK_DEFERRED_ERROR

Severity	warning
Description	Deferred error on <i>Disk ID</i> , start LBA= <i>Start LBA</i> , last LBA= <i>Last LBA</i> , latency= <i>latency</i> msec, key= <i>key</i>
Troubleshooting	N/A

SSD_DEFERRED_ERROR

Severity	warning
Description	SSD <i>SSD ID</i> signaled deferred error on start lba= <i>first_lba</i> , last lba= <i>last_lba</i> , scsi_opcode= <i>scsi_opcode</i> , latency= <i>latency</i> usec, key= <i>key</i>
Troubleshooting	N/A

DISK_MEDIUM_ERROR

Severity	warning
Description	Media errors on <i>Disk ID</i> , start LBA= <i>Start LBA</i> , last LBA= <i>Last LBA</i> , latency= <i>latency</i> msec.
Troubleshooting	N/A

SSD_MEDIUM_ERROR

Severity	warning
Description	Media errors on <i>SSD ID</i> , start LBA= <i>Start LBA</i> , last LBA= <i>Last LBA</i> , latency= <i>latency</i> msec.
Troubleshooting	N/A

DISK_ABNORMAL_ERROR

Severity	major
Description	Unit attentions or aborts in the last 30 minutes on <i>Disk ID</i> , start lba= <i>start_lba</i> , last lba= <i>last_lba</i> , command= <i>command</i> , latency= <i>latency</i> msec.
Troubleshooting	N/A

SSD_ABNORMAL_ERROR

Severity	major
Description	Unit attentions or aborts in the last 30 minutes on <i>SSD ID</i> , start lba= <i>start_lba</i> , last lba= <i>last_lba</i> , command= <i>command</i> , latency= <i>latency</i> msec.
Troubleshooting	N/A

DISK_LONG_LATENCY

Severity	variable
Description	Long latencies on disk I/Os in the last 30 minutes on <i>Disk ID</i> , start LBA= <i>Start LBA</i> , last LBA= <i>Last LBA</i> , command= <i>command</i> , latency= <i>latency</i> msec.
Troubleshooting	N/A

SSD_LONG_LATENCY

Severity	variable
Description	Long latencies on ssd I/Os in the last 30 minutes on <i>SSD ID</i> , start LBA= <i>Start LBA</i> , last LBA= <i>Last LBA</i> , scsi_opcode= <i>scsi_opcode</i> , latency= <i>latency</i> msec.

DISK_BAD_PERFORMANCE

Severity	minor
Description	Bad performance on <i>Disk ID</i> , I/O count= <i>I/O Count</i> , transferred kbytes= <i>kbytes</i> ,msecs= <i>seconds</i> .
Troubleshooting	N/A

SSD_BAD_PERFORMANCE

Severity	major
Description	Bad performance on <i>SSD ID</i> , I/O count= <i>I/O Count</i> , transferred kbytes= <i>kbytes</i> ,msecs= <i>seconds</i> .
Troubleshooting	N/A

DISK_UNRESPONSIVE

Severity	major
Description	Disk <i>Disk ID</i> is unresponsive for <i>time</i> msecs, cache dirty level is <i>Dirty Level</i> %

SSD_UNRESPONSIVE

Severity	major
Description	SSD <i>SSD ID</i> is unresponsive for <i>time</i> msecs

DISK_RESPONSIVE

Severity	major
Description	Disk <i>Disk ID</i> is now responsive. Was unresponsive for <i>unresponsive_time</i> msecs, cache dirty level is <i>Dirty Level</i> %

SSD_RESPONSIVE

Severity	major
Description	SSD <i>SSD ID</i> is now responsive. Was unresponsive for <i>time</i> msecs

MIRROR_CREATE

Severity	informational
Description	A remote mirror was defined for Volume ' <i>local volume name</i> ' on Target ' <i>target name</i> '. Remote Volume is ' <i>remote volume name</i> '.

CG_MIRROR_CREATE

Severity	informational
Description	A remote mirror was defined for Consistency Group ' <i>local CG name</i> ' on Target ' <i>target name</i> '. Remote Consistency Group is ' <i>remote CG name</i> '.

MIRROR_CREATE_SLAVE

Severity	informational
Description	A remote mirror was defined by Target ' <i>target name</i> ' for Volume ' <i>local volume name</i> '. Remote Volume is ' <i>remote volume name</i> '.

CG_MIRROR_CREATE_SLAVE

Severity	informational
Description	A remote mirror was defined by Target ' <i>target name</i> ' for CG ' <i>local CG name</i> '. Remote CG is ' <i>remote CG name</i> '.

MIRROR_SCHEDULE_CHANGE

Severity	informational
Description	Schedule of remote mirror of ' <i>local peer name</i> ' is now ' <i>schedule name</i> '.

MIRROR_CREATE_FAILED_TARGET_NOT_CONNECTED

Severity	warning
Description	Target could not be reached. Target with name ' <i>target.name</i> ' is currently not connected.
Troubleshooting	Connect the target system to this system.

REMOTE_OPERATION_FAILED_TIMED_OUT

Severity	warning
Description	Operation on remote machine timed out. Invoking ' <i>Function Name</i> ' on target ' <i>Target Name</i> ' timed out.
Troubleshooting	Retry operation. If problem persists contact support.

MIRROR_RESYNC_FAILED

Severity	major
Description	Synchronization of meta data with mirror failed. Configuration of remote mirror of volume ' <i>local volume name</i> ' on target ' <i>target name</i> ' does not match local configuration.
Troubleshooting	Make sure configuration on both machines is compatible and activate the mirror. If problem persists contact support.

MIRROR_RESYNC_FAILED_DUE_TO_THIN_PROVISIONING

Severity	major
Description	Synchronization of bitmaps with mirror failed. Not enough hard capacity left in Pool of volume ' <i>mirror.local_volume_name</i> '.
Troubleshooting	Delete unnecessary volumes in pool or enlarge the pool's hard size.

MIRROR_SYNC_STARTED

Severity	informational
Description	Synchronization of remote mirror of volume ' <i>local volume name</i> ' on Target ' <i>target name</i> ' has started.

MIRROR_SYNC_ENDED

Severity	informational
Description	Synchronization of remote mirror of peer ' <i>local peer name</i> ' on target ' <i>target name</i> ' has ended.

MIRROR_CANNOT_CREATE_SYNC_JOB_TOO_MANY_VOLUMES

Severity	major
Description	Synchronization of remote mirror of peer ' <i>local peer name</i> ' on target ' <i>target name</i> ' can not be synced , insufficient volume available for this operation.

MIRROR_CANNOT_CREATE_LRS_TOO_MANY_VOLUMES

Severity	major
Description	Synchronization of remote mirror of peer ' <i>local peer name</i> ' on target ' <i>target name</i> ' can not be synced , insufficient volume available for this operation.

MIRROR_REESTABLISH_FAILED

Severity	major
Description	Mirror reestablish failed. Connection to remote mirror of peer ' <i>local peer name</i> ' on target ' <i>target name</i> ' could not be established. Remote action failed.

MIRROR_REESTABLISH_FAILED_CONFIGURATION_ERROR

Severity	major
Description	Mirror reestablish failed. Either configuration of remote mirror of peer ' <i>local peer name</i> ' on target ' <i>target name</i> ' does not match local configuration.
Troubleshooting	Make sure configuration on both machines is compatible and activate the mirror. If problem persists contact support.

MIRROR_ACTIVATE

Severity	informational
Description	The Remote Mirror of peer ' <i>local peer name</i> ' on Target ' <i>target name</i> ' was activated.

MIRROR_DEACTIVATE

Severity	informational
Description	The Remote Mirror of peer ' <i>local peer name</i> ' on Target ' <i>target name</i> ' was deactivated.

MIRROR_DEACTIVATE_SECONDARY_LOCKED

Severity	minor
Description	The Remote Mirror of peer ' <i>local peer name</i> ' on Target ' <i>target name</i> ' was deactivated since the Pool on the secondary machine was locked.

MIRROR_DEACTIVATE_CONFIGURATION_ERROR

Severity	minor
Description	The Remote Mirror of peer ' <i>local peer name</i> ' on Target ' <i>target name</i> ' was deactivated since the Mirror configuration on the slave machine has changed.

MIRROR_DELETE

Severity	informational
Description	The Remote Mirror relation of peer ' <i>local peer name</i> ' to a peer on Target ' <i>target name</i> ' was deleted.

MIRROR_REVERSE_ROLE_TO_SLAVE

Severity	informational
Description	Local peer ' <i>local peer name</i> ' is now Slave of a peer on Target ' <i>target name</i> '.

MIRROR_REVERSE_ROLE_TO_MASTER

Severity	informational
Description	Local peer ' <i>local peer name</i> ' is now Master of a peer on Target ' <i>target name</i> '.

MIRROR_REVERSE_ROLE_OF_PEER_WITH_LCS_TO_MASTER

Severity	informational
Description	Local peer ' <i>local peer name</i> ' is now Master of a peer on Target ' <i>target name</i> ' external last consistent snapshot should be deleted manually .

MIRROR_SWITCH_ROLES_TO_SLAVE

Severity	informational
Description	Local peer ' <i>local peer name</i> ' switched roles with peer on Target ' <i>target name</i> '. It is now Slave.

MIRROR_SWITCH_ROLES_TO_MASTER

Severity	informational
Description	Local peer ' <i>local peer name</i> ' switched roles with peer on Target ' <i>target name</i> '. It is now Master.

MIRROR_REESTABLISH_FAILED_TOO_MANY_VOLUMES

Severity	major
Description	Last Consistent Snapshot of Slave peer ' <i>local peer name</i> ' could not be created. Maximal number of Volumes are already defined.
Troubleshooting	Delete Volumes to allow new ones to be created. Activate Mirror on the Master Machine.

MIRROR_END_SYNC_FAILED_CONFIGURATION_ERROR

Severity	major
Description	Configuration of remote mirror of peer ' <i>local peer name</i> ' on target ' <i>target name</i> ' does not match local configuration.
Troubleshooting	Make sure configuration on both machines is compatible and activate the mirror. If problem persists contact support.

MIRROR_CHANGE_DESIGNATION

Severity	informational
Description	Local peer ' <i>local peer name</i> ' switched its designated role with peer on Target ' <i>target name</i> '. It is now <i>designation</i> .

MIRROR_CANCEL_SNAPSHOT

Severity	informational
Description	All mirrored snapshots which were created for Mirror of peer ' <i>local peer name</i> ' and were not yet synchronized will not be mirrored in the remote machine.

MIRROR_SYNCHRONIZATION_TYPE_CHANGED

Severity	informational
Description	Synchronization of Mirror of peer ' <i>local peer name</i> ' is now ' <i>mirror synchronization type</i> '.

DM_DEFINE

Severity	informational
Description	Data Migration was defined to Volume ' <i>local volume name</i> ' from Target ' <i>target name</i> '.

OM_DEFINE

Severity	informational
Description	Online Migration was defined to Volume ' <i>local volume name</i> ' from Target ' <i>target name</i> '.

DM_SYNC_STARTED

Severity	informational
Description	Migration to volume ' <i>local volume name</i> ' from Target ' <i>target name</i> ' has started.

DM_SYNC_ENDED

Severity	informational
Description	Migration to volume ' <i>local volume name</i> ' from target ' <i>target name</i> ' is complete.

DM_SYNC_ENDED_WITH_ERRORS

Severity	Critical
Description	Migration to volume ' <i>local volume name</i> ' from target ' <i>target name</i> ' has completed with <i>medium_errors_in_data_migration</i> error(s). Check previous events related to this volume for the list of affected LBAs.'.

DM_ACTIVATE

Severity	informational
Description	Migration to Volume ' <i>local volume name</i> ' from Target ' <i>target name</i> ' was activated.

DM_DEACTIVATE

Severity	informational
Description	Migration to Volume ' <i>local volume name</i> ' from Target ' <i>target name</i> ' was deactivated.

DM_DEACTIVATE_LUN_UNAVAILABLE

Severity	minor
Description	Migration to Volume ' <i>local volume name</i> ' from Target ' <i>target name</i> ' was deactivated since LUN is not available on one of the active paths to the target.

DM_START_MIGRATION

Severity	informational
Description	Migration to Volume ' <i>local volume name</i> ' from Target ' <i>target name</i> ' will now start automatically.

DM_DELETE

Severity	informational
Description	Definition of Data Migration to Volume ' <i>local volume name</i> ' from Target ' <i>target name</i> ' was deleted.

SCHEDULE_CREATE

Severity	informational
Description	Schedule was created with name ' <i>schedule name</i> '.

SCHEDULE_UPDATE

Severity	informational
Description	Schedule with name ' <i>schedule name</i> ' was updated.

SCHEDULE_RENAME

Severity	informational
Description	Schedule with name ' <i>old_name</i> ' was renamed ' <i>schedule name</i> '.

SCHEDULE_DELETE

Severity	informational
Description	Schedule with name ' <i>schedule name</i> ' was deleted.

MIRROR_RPO_OK

Severity	informational
Description	Mirror of local peer ' <i>local peer name</i> ' is now ahead of its specified RPO.

MIRROR_RPO_LAGGING

Severity	informational
Description	Mirror of local peer ' <i>local peer name</i> ' is now behind its specified RPO.

MIRROR_CHANGE_RPO

Severity	informational
Description	RPO or Mirror of local peer ' <i>local peer name</i> ' is now <i>RPO</i> .

MIRROR_IS_LAGGING_BEYOND_PERCENT_THRESHOLD

Severity	warning
Description	Last Replication Time of Mirror of local peer ' <i>local peer name</i> ' is <i>Last Replication Time</i> .

MIRROR_AUTO_FIX_REACHED_LIMIT

Severity	warning
Description	A remote checksum diff for mirror ' <i>local peer name</i> ' cannot be fixed automatically because we reached the auto fix limit.

MIRROR_IS_LAGGING_BEYOND_ABSOLUTE_THRESHOLD

Severity	warning
Description	Last Replication Time of Mirror of local peer ' <i>local peer name</i> ' is <i>Last Replication Time</i> .

MIRROR_INCOMPATIBLE_VERSION_FOR_UNMAP_SUPPORT

Severity	warning
Description	Mirror of peer ' <i>local peer name</i> ' on target ' <i>target name</i> ' cannot support unmap, remote machine has incompatible version.

XMIRROR_DEFINE

Severity	informational
Description	A xmirror master ' <i>xmirror name</i> ' was defined for volume ' <i>local volume name</i> '.

XMIRROR_DEFINE_SLAVE

Severity	informational
Description	A xmirror slave ' <i>xmirror name</i> ' was defined for volume ' <i>local volume name</i> '.

XMIRROR_DEFINE_SMASTER

Severity	informational
Description	A xmirror smaster ' <i>xmirror name</i> ' was defined for volume ' <i>local volume name</i> '.

XMIRROR_ACTIVATE

Severity	informational
Description	Xmirror ' <i>xmirror name</i> ' was activated.

XMIRROR_DEACTIVATE

Severity	informational
Description	Xmirror ' <i>xmirror name</i> ' was deactivated.

XMIRROR_DELETE

Severity	informational
Description	Xmirror ' <i>xmirror name</i> ' was deleted.

XMIRROR_CHANGE_SLAVE_ROLE_TO_MASTER

Severity	informational
Description	Xmirror ' <i>xmirror name</i> ' was changed to standalone.

XMIRROR_CHANGE_MASTER_ROLE_TO_SLAVE

Severity	informational
Description	Xmirror ' <i>xmirror name</i> ' was changed to slave.

XMIRROR_CHANGE_SMASTER_ROLE_TO_MASTER

Severity	informational
Description	Xmirror ' <i>xmirror name</i> ' was changed to master.

XMIRROR_CHANGE_MASTER_ROLE_TO_SMASTER

Severity	informational
Description	Xmirror ' <i>xmirror name</i> ' was changed to smaster.

XMIRROR_STANDBY_MIRROR_REGISTERED

Severity	informational
Description	Xmirror ' <i>xmirror name</i> ' registered a standby mirror on SMASTER system

XMIRROR_COMPROMISED

Severity	warning
Description	Xmirror ' <i>xmirror name</i> ' is compromised, reason: <i>Compromise Reason</i>

XMIRROR_RESTORED

Severity	informational
Description	Xmirror ' <i>xmirror name</i> ' restored after being compromised

XMIRROR_RENAMED

Severity	informational
Description	Xmirror ' <i>Old Xmirror Name</i> ' was renamed to ' <i>xmirror name</i> '.

MAP_VOLUME

Severity	informational
Description	Volume with name ' <i>volume.name</i> ' was mapped to LUN ' <i>LUN</i> ' for <i>host_or_cluster</i> with name ' <i>host</i> '.

MAP_PROXY_VOLUME

Severity	informational
Description	IBM Hyper-Scale Mobility Volume with name ' <i>name</i> ' was mapped to LUN ' <i>LUN</i> ' for <i>host_or_cluster</i> with name ' <i>host</i> '.

UNMAP_VOLUME

Severity	informational
Description	Volume with name ' <i>volume.name</i> ' was unmapped from <i>host_or_cluster</i> with name ' <i>host</i> '.

UNMAP_PROXY_VOLUME

Severity	informational
Description	IBM Hyper-Scale Mobility Volume with name ' <i>name</i> ' was unmapped from <i>host_or_cluster</i> with name ' <i>host</i> '.

CLUSTER_DEFINE_EXCEPTION

Severity	informational
Description	LUN ' <i>LUN</i> ' was defined as having host specific mapping in cluster ' <i>cluster</i> '.

CLUSTER_CANCEL_EXCEPTION

Severity	informational
Description	LUN ' <i>LUN</i> ' was defined as having uniform mapping in cluster ' <i>cluster</i> '.

SPECIAL_TYPE_SET

Severity	informational
Description	Type of <i>host_or_cluster</i> with name ' <i>host</i> ' was set to ' <i>type</i> '.

DATA_LOSS

Severity	critical
Description	Manager found data loss in <i>num_lost_slices</i> slices.
Troubleshooting	Please contact support.

SERVICE_HAS_FAILED

Severity	major
Description	<i>Component ID</i> has failed.
Troubleshooting	Please contact support.

SERVICE_FAILED_TO_PHASEIN

Severity	major
Description	<i>Component ID</i> failed to phase-in.
Troubleshooting	Please contact support.

SERVICE_FAILED_TO_RESTART

Severity	major
Description	<i>Component ID</i> failed to restart.
Troubleshooting	Please contact support.

MANAGER_RESIGNED_TO_LET_MODULE_PHASE_OUT

Severity	informational
Description	Previous manager resigned to let <i>Component ID</i> phase out.
Troubleshooting	N/A

MODULE_FAILED

Severity	critical
Description	<i>Component ID</i> failed.
Troubleshooting	Please contact support.

NODE_FAILED

Severity	critical

Description	Node # <i>Node ID</i> of type <i>Node Type</i> on <i>Module Component ID</i> failed because of <i>failure_reason</i> .
Troubleshooting	Please contact support.

NODE_RESET

Severity	informational
Description	Node # <i>Node ID</i> of type <i>Node Type</i> on <i>Module Component ID</i> was reset.
Troubleshooting	N/A

NODE_IS_NOT_UP

Severity	minor
Description	Node # <i>Node ID</i> of type <i>Node Type</i> on <i>Module Component ID</i> is not up.
Troubleshooting	Please contact support

TRACES_SNAPSHOT_CREATE

Severity	informational
Description	Traces snapshot created. ID <i>snap_id</i>
Troubleshooting	N/A

MODULE_IS_NOT_UP

Severity	minor
Description	<i>Module Component ID</i> is not up.
Troubleshooting	Please contact support

SINGLETON_NODE_IS_NOT_UP

Severity	minor
Description	No singleton of type <i>Node Type</i> is up.
Troubleshooting	Please contact support

DISK_HAS_FAILED

Severity	variable
Description	Disk <i>Component ID</i> Failed.
Troubleshooting	Please contact your Administrator.

SSD_HAS_FAILED

Severity	major
Description	SSD <i>Component ID</i> Failed.
Troubleshooting	Please contact your Administrator.

DISK_INFO_EXTRA_EVENT

Severity	informational
Description	Disk <i>Component ID</i> extra information event.
Troubleshooting	N/A

SSD_INFO_EXTRA_EVENT

Severity	informational
Description	SSD <i>Component ID</i> extra information event.
Troubleshooting	N/A

COMPONENT_TEST_OF_DISK_HAS_FAILED

Severity	major
Description	Test of <i>Component ID</i> has failed with error <i>Error</i> .
Troubleshooting	Please contact your Administrator.

COMPONENT_TEST_OF_SSD_HAS_FAILED

Severity	major
Description	Test of <i>Component ID</i> has failed with error <i>Error</i> .
Troubleshooting	Please contact your Administrator.

DISK_INFO_LOAD_FAILED

Severity	major
Description	<i>Component ID</i> failed.
Troubleshooting	Please contact your Administrator.

DISK_STARTED_PHASEOUT

Severity	informational
Description	System started phasing out <i>Component ID</i> .
Troubleshooting	N/A

DISK_STARTED_AUTO_PHASEOUT

Severity	minor
Description	System started automatic phasing out <i>Component ID</i> .
Troubleshooting	Please contact support.

DISK_STARTED_PHASEIN

Severity	informational
Description	System started phasing in <i>Component ID</i> .
Troubleshooting	N/A

DISK_FINISHED_PHASEIN

Severity	informational
Description	System finished phasing in <i>Component ID</i> .
Troubleshooting	N/A

DISK_FINISHED_PHASEOUT

Severity	informational
Description	System finished phasing out <i>Component ID</i> .
Troubleshooting	N/A

SECOND_DISK_FAILURE

Severity	critical
Description	Disk <i>Component ID</i> failed during rebuild.
Troubleshooting	N/A

DISK_RECOVERED

Severity	critical
Description	Disk <i>Component ID</i> is functioning again.
Troubleshooting	N/A

MODULE_STARTED_PHASEOUT

Severity	informational
Description	System started phasing out <i>Component ID</i> .
Troubleshooting	N/A

MODULE_INFO_PRE_EVENT

Severity	informational
Description	<i>Component ID</i> information pre event.
Troubleshooting	N/A

MODULE_FINISHED_PHASEOUT

Severity	informational
Description	System finished phasing out <i>Component ID</i> .
Troubleshooting	N/A

MODULE_STOPPED_PHASEOUT_DUE_TO_MANAGEMENT_REQUIREMENT

Severity	major
Description	System stopped phasing out <i>Component ID</i> due to management requirement.
Troubleshooting	N/A

MODULE_PHASEOUT_FAILURE_REASON

Severity	informational
Description	System could not phaseout <i>Component ID</i> due to lack of nodes of type <i>Node Type</i> .
Troubleshooting	N/A

START_WORK

Severity	informational
Description	System has entered ON state.

SYSTEM_HAS_ENTERED_MAINTENANCE_MODE

Severity	warning
Description	System has entered MAINTENANCE state [<i>Reason</i>]

SYSTEM_ENTERED_CHARGING_STATE

Severity	informational
Description	System cannot start work until it is sufficiently charged.

SYSTEM_LEFT_CHARGING_STATE

Severity	informational
Description	System is sufficiently charged.

USER_SHUTDOWN

Severity	major
Description	System is shutting down due to a user request.

EMERGENCY_SHUTDOWN_NOW

Severity	critical
Description	System is shutting down in emergency shutdown mode due to: <i>Emergency Shutdown Reason</i> .
Troubleshooting	Please contact your Administrator.

SHUTDOWN_PARAMS

Severity	major
Description	System action is ' <i>Shutdown Action</i> '. Target state is ' <i>Target State</i> '. Safemode is ' <i>Safe Mode</i> '. UPS Sleep Time= <i>UPS sleep time in seconds</i> seconds.

DISK_STARTED_AUTO_PHASEIN

Severity	critical
Description	System started phasing in <i>Component ID</i> in order to ensure that data will not be unprotected. Phaseout of the containing service and module has been cancelled.
Troubleshooting	N/A

SANITY_CHECK_FAILED

Severity	critical

Description	Sanity check failed.
Troubleshooting	Please contact support.

SYSTEM_HARD_CAPACITY_CHANGED

Severity	informational
Description	System's hard capacity is now <i>Capacity</i> GB.

SYSTEM_DISK_CAPACITY_EXPANDED

Severity	informational
Description	System's hard capacity is now <i>Capacity</i> GB.

SYSTEM_DISK_CAPACITY_LIMIT_PERCENTAGE_EXPANDED

Severity	informational
Description	System's hard capacity was expanded to <i>Capacity limit Percentage</i> .

SYSTEM_CAN_NOT_INCREASE_SPARES

Severity	informational
Description	System's spares can not be increased to <i>modules</i> modules and <i>disks</i> disks. <i>Capacity</i> GB should be freed.

SYSTEM_SOFT_CAPACITY_CHANGED

Severity	informational
Description	System's soft capacity is now <i>Capacity</i> GB.

MODULE_IS_MISSING_DATA_DISKS

Severity	major
Description	<i>Module ID</i> has <i>Num Found</i> of <i>Num Expected</i> data disks.

SERVICE_WAS_RESTARTED

Severity	informational
Description	<i>Module ID</i> was restarted.

DATA_SERVICE_STARTED_PHASEOUT

Severity	informational
Description	System started phasing out <i>Component ID</i> .
Troubleshooting	N/A

DATA_SERVICE_FINISHED_PHASEOUT

Severity	informational
Description	System finished phasing out <i>Component ID</i> .
Troubleshooting	N/A

POWER_SUPPLY_UNIT_STATUS_IS_OK

Severity	informational
Description	The status of <i>Component ID</i> is now OK.
Troubleshooting	Please contact support.

POWER_TO_MODULE_SHOULD_BE_DISCONNECTED_AND_RECONNECTED

Severity	informational
Description	<i>Component ID</i> should be disconnected completely from power supply and reconnected after 60 seconds in order to complete the firmware upgrade scheme.
Troubleshooting	N/A

MODULE_FIRMWARE_UPGRADE_FAILED

Severity	warning
Description	<i>Module Component ID</i> has failed due to <i>failure_reason</i> .
Troubleshooting	Please contact support

FC_CHANGE_DETECTED

Severity	major
Description	<i>Component ID</i> has been changed from a <i>Old Model</i> with a serial of <i>old_serial</i> to a <i>New Model</i> with a serial of <i>new_serial</i> .
Troubleshooting	Was this fiber channel port actually replaced?

DISK_SMALLER_THAN_SYSTEM_DISK_SIZE

Severity	major
Description	Disk <i>Component ID</i> has a size of <i>New sizeGB</i> which is smaller than system disk size <i>System sizeGB</i> .

SSD_SMALLER_THAN_SYSTEM_SSD_SIZE

Severity	informational
Description	SSD <i>Component ID</i> has a size of <i>New sizeGB</i> which is smaller than system ssd size <i>System sizeGB</i> .

DISK_LARGER_THAN_SYSTEM_DISK_SIZE

Severity	variable
Description	Disk <i>Component ID</i> has a size of <i>New sizeGB</i> which is larger than system disk size <i>System sizeGB</i> .

SSD_LARGER_THAN_SYSTEM_SSD_SIZE

Severity	major
Description	SSD <i>Component ID</i> has a size of <i>New sizeGB</i> which is larger than system ssd size <i>System sizeGB</i> .

DISK_MODEL_DIFFERENT_THAN_SYSTEM_MODEL

Severity	major
Description	Disk <i>Component ID</i> of model <i>Disk model</i> , by vendor <i>Disk vendor</i> , differs from system disk by vendor <i>System vendor</i> and model <i>System model</i> .

DISK_MODEL_DIFFERENT_THAN_MODULE_DISK_MODEL

Severity	major
Description	Disk <i>Component ID</i> of model <i>Disk model</i> , by vendor <i>Disk vendor</i> , differs from module's disk model <i>Module model</i> , by vendor <i>Module vendor</i> .

FIRST_DISK_MODEL_IN_MODULE

Severity	informational
Description	Disk <i>Component ID</i> of model <i>Disk model</i> , by vendor <i>Disk vendor</i> , was the first to be added to a module with an inter-module disk intermix policy, this will require all other disks to have the same model and vendor.

SSD_INTERMIX_DETECTED

Severity	informational
Description	SSD <i>Component ID</i> of model <i>SSD model</i> , by vendor <i>SSD vendor</i> , <i>User message Required model</i>

SSD_CACHING_ENABLED

Severity	informational
Description	SSD Caching feature enabled. SSDs can now be installed.
Troubleshooting	N/A

SSD_CACHING_DISABLED

Severity	informational
Description	SSD Caching feature disabled.
Troubleshooting	N/A

FLASH_CACHE_ENABLE

Severity	informational
Description	Flash Cache feature enabled. SSDs can now be installed.
Troubleshooting	N/A

FLASH_CACHE_DISABLE

Severity	informational
Description	Flash Cache feature disabled.
Troubleshooting	N/A

ENCRYPT_ENABLE_STARTED

Severity	informational
Description	Starting disk encryption activation. This process can take several minutes to complete.

ENCRYPT_ENABLE_COMPLETED

Severity	informational
Description	Disk encryption is in effect.

ENCRYPT_ENABLE_NOT_COMPLETED

Severity	major
Description	Cannot complete encryption activation. <i>Count</i> disk(s) could not be enrolled.
Troubleshooting	Please contact technical support

ENCRYPT_DISABLE_STARTED

Severity	informational
Description	Starting disk encryption deactivation process.

ENCRYPT_DISABLE_COMPLETED

Severity	informational
Description	Disk encryption is no longer in effect.

ENCRYPT_DISABLE_NOT_COMPLETED

Severity	major
Description	Cannot complete encryption deactivation. <i>Count</i> disk(s) could not be securely erased.
Troubleshooting	Please contact technical support

ENCRYPT_KEYSERVER_ADDED

Severity	informational
Description	A key server named ' <i>Key Server Name</i> ' was added.

ENCRYPT_KEYSERVER_DELETED

Severity	informational
Description	Key server ' <i>Key Server Name</i> ' was deleted.

ENCRYPT_KEYSERVER_EDITED

Severity	informational
Description	Details of key server ' <i>Key Server Name</i> ' were modified.

ENCRYPT_KEYSERVER_RENAMED

Severity	informational
Description	Key server ' <i>Old Name</i> ' was renamed to ' <i>New Name</i> '.

ENCRYPT_KEYSERVER_REKEY_COMPLETED

Severity	informational
Description	Key server ' <i>Key Server Name</i> ' rekey completed.

ENCRYPT_UNABLE_TO_UPDATE_KEY_DURING_DEACTIVATE_ON_KEYSERVER

Severity	major
Description	Could not update key server ' <i>Keyserver Name</i> ' regarding encryption deactivation. Please check key server status.

ENCRYPT_UNABLE_TO_UPDATE_KEY_DURING_DEACTIVCATE_ON_KEYSERVER

Severity	major
Description	Could not update key server ' <i>Keyserver Name</i> ' regarding encryption deactivation. Please check key server status.

ENCRYPT_KEYSERVER_REKEY_FAILED

Severity	major
Description	Cannot complete rekey with key server ' <i>Key Server Name</i> '.

ENCRYPT_KEYSERVER_REKEY_ROLLBACK_FAILED

Severity	major
Description	Cannot rollback failed rekey with key server ' <i>Key Server Name</i> '.

ENCRYPT_RECOVERY_KEY_ENTERED

Severity	informational
Description	Valid recovery key share was entered by user ' <i>User Name</i> '.

ENCRYPT_INVALID_RECOVERY_KEY_ENTERED

Severity	major
Description	Invalid recovery key share was entered by user ' <i>User Name</i> '.

ENCRYPT_RECOVERY_KEYS_GENERATED

Severity	informational
Description	Recovery keys created.

ENCRYPT_RECOVERY_KEY_REKEY_SUCCESS

Severity	informational
Description	Recovery key rekey was successful.

ENCRYPT_RECOVERY_KEY_REKEY_FAIL

Severity	major
Description	Recovery key rekey failed.

ENCRYPT_RECOVERY_KEY_VERIFIED

Severity	informational
Description	Recovery key verified successfully for user ' <i>User Name</i> '.

ENCRYPT_RECOVERY_KEY_VERIFY_FAILED

Severity	major
Description	Recovery key verification failed for user ' <i>User Name</i> '.

ENCRYPT_RECOVERY_KEY_ALL_SHARES_VERIFIED

Severity	informational
Description	All recovery key shares have been verified.

ENCRYPT_KR_WRITE_FAILED

Severity	critical
Description	Key repository write failed with error code <i>rc</i> .
Troubleshooting	Please contact technical support

ENCRYPT_KR_READ_FAILED

Severity	major
Description	Key repository read failed with error code <i>rc</i> .
Troubleshooting	Please contact technical support

ENCRYPT_KEYSERVER_UNREACHABLE

Severity	minor
Description	Cannot ping key server ' <i>Keyserver Name</i> ' <i>TEXT</i> from module <i>module</i> .

ENCRYPT_UNABLE_TO_RETRIEVE_KEY_FROM_KEYSERVER

Severity	major
Description	Failed to retrieve key from key server ' <i>Keyserver Name</i> ' via <i>TEXT</i> on module <i>node id</i> . Please verify that the key server type and version are supported. If so, please check its status.

KEYSERVER_GATEWAY_FAILED_TO_REPORT_STATUS

Severity	minor
Description	Key server gateway in module <i>module</i> failed to report key server's status: <i>message</i> .
Troubleshooting	Please contact technical support

ENCRYPT_RECOVERY_KEY_RECOVER_SUCCESSFUL

Severity	informational
Description	Key recovery was successful, unlocking system.

ENCRYPTION_CERTIFICATE_FOR_XIV_IS_NOT_INSTALLED

Severity	critical
Description	XIV certificate is not installed.
Troubleshooting	Check output of <i>pki_list</i> for a certificate named XIV and contact technical support

ENCRYPT_UNABLE_TO_DELETE_MASTER_KEYSERVER

Severity	informational
Description	Deletion of master key server ' <i>Keyserver Name</i> ' is not allowed. Please define another key server as master first.

ENCRYPTION_SKMIP_ERROR

Severity	major
Description	Module <i>Module</i> reported <i>Keyserver Name</i> returned error: <i>error code</i> - <i>TEXT</i>
Troubleshooting	Please contact the next level of support.

SSD_HAS_FAILED_WHILE_ENABLING_ENCRYPTION

Severity	major
Description	SSD <i>Component ID</i> failed while enabling encryption. User data on the device may not have been cryptographically erased. The failed device should be handled accordingly.
Troubleshooting	Contact support

CF_FAILED

Severity	major
Description	<i>Component ID</i> has failed. Hardware status: <i>Status</i> .
Troubleshooting	Please contact support.

DIMM_FAILED

Severity	major
Description	<i>Component ID</i> has failed. Hardware status: <i>Status</i> .
Troubleshooting	Please contact support.

CPU_FAILED

Severity	major
Description	<i>Component ID</i> has failed. Hardware status: <i>Status</i> .
Troubleshooting	Please contact support.

NIC_FAILED

Severity	major
Description	<i>Component ID</i> has failed. Hardware status: <i>Status</i> .
Troubleshooting	Please contact support.

UNKNOWN_MODULE_SERIAL_NUMBER

Severity	critical
Description	<i>Component ID</i> has an unknown serial number of <i>serial</i> .
Troubleshooting	Who generates this serial number?

ILLEGAL_MODULE_SERIAL_NUMBER

Severity	critical
Description	<i>Component ID</i> has an illegal serial number of <i>serial</i> .
Troubleshooting	Is this a Pre-GA module?

MODULE_ORIGINAL_SERIAL_CLEANED

Severity	major
Description	<i>Component ID</i> had an illegal original serial number of <i>Original Serial</i> which was changed to <i>Clean Original Serial</i> .
Troubleshooting	Is this a Pre-GA module?

MODULE_ORIGINAL_PART_NUMBER_CLEANED

Severity	major
Description	<i>Component ID</i> had an illegal original part number of <i>Original Serial</i> which was changed to <i>Clean Original Serial</i> .
Troubleshooting	Is this a Pre-GA module?

ILLEGAL_MODULE_PART_NUMBER

Severity	critical
Description	<i>Component ID</i> has an illegal part number of <i>part_number</i> .
Troubleshooting	Is there a problem with the IPMI of the module?

MODULE_CHANGE_DETECTED

Severity	major
Description	<i>Component ID</i> has been changed from a serial of <i>old_serial</i> to <i>new_serial</i> .
Troubleshooting	Was this module actually replaced?

CF_CHANGE_DETECTED

Severity	major
Description	<i>Component ID</i> has had the CF changed from a serial of <i>old_serial</i> to <i>new_serial</i> .
Troubleshooting	Was this CF actually replaced?

INFINIBAND_HCA_CHANGE_DETECTED

Severity	major
Description	Adapter Module serial number changed from <i>old_serial</i> to <i>new_serial</i> .

Troubleshooting	Was this InfiniBand HCA actually replaced?

INFINIBAND_HCA_BOARD_CHANGE_DETECTED

Severity	major
Description	Adapter <i>Module</i> board type changed from ' <i>old_board_description</i> ' to ' <i>new_board_description</i> '.
Troubleshooting	Was this InfiniBand HCA actually replaced?

INFINIBAND_HCA_BOARD_ID_CHANGE_DETECTED

Severity	major
Description	Adapter <i>Module</i> board ID changed from ' <i>old board id</i> ' to ' <i>new board id</i> '.
Troubleshooting	Was this InfiniBand HCA actually replaced?

CNA_CHANGE_DETECTED

Severity	major
Description	Adapter <i>CNA Id</i> serial number changed from <i>old_serial</i> to <i>new_serial</i> .
Troubleshooting	Was this CNA actually replaced?

CNA_BOARD_CHANGE_DETECTED

Severity	major
Description	Adapter <i>CNA Id</i> board type changed from ' <i>old_board_description</i> ' to ' <i>new_board_description</i> '.
Troubleshooting	Was this CNA actually replaced?

CNA_BOARD_ID_CHANGE_DETECTED

Severity	major
Description	Adapter <i>CNA Id</i> board ID changed from ' <i>old board id</i> ' to ' <i>new board id</i> '.
Troubleshooting	Was this CNA actually replaced?

INFINIBAND_HCA_ADAPTER_TYPE_CHANGE_DETECTED

Severity	major
Description	Adapter <i>HCA Id</i> adapter type changed from ' <i>old_board_type</i> ' to ' <i>new_board_type</i> '.
Troubleshooting	Was this adapter actually replaced?

SAS_CONTROLLER_CHANGE_DETECTED

Severity	major
Description	The SAS controller on module <i>Module</i> was changed from a serial of <i>old_serial</i> and board assembly of ' <i>old_assembly</i> ' to serial <i>new_serial</i> and board assembly ' <i>new_assembly</i> '.
Troubleshooting	Was this SAS controller actually replaced?

DIMM_CHANGE_DETECTED

Severity	major
Description	<i>Component ID</i> has been changed from a serial of <i>old_serial</i> to <i>new_serial</i> .
Troubleshooting	Was this DIMM actually replaced?

PSU_CHANGE_DETECTED

Severity	informational
Description	<i>Component ID</i> has been changed from a serial number ' <i>old_serial</i> ', part number ' <i>old_part_number</i> ', to serial number ' <i>new_serial</i> ' and part number ' <i>new_part_number</i> '.
Troubleshooting	Was this PSU actually replaced?

PSU_WAS_REMOVED

Severity	warning
Description	<i>Component ID</i> with a serial number ' <i>Serial</i> ' and part number ' <i>Part Number</i> ' was removed from the system.
Troubleshooting	Was this PSU actually removed?

PSU_WAS_INSTALLED

Severity	informational
Description	<i>Component ID</i> with a serial number ' <i>Serial</i> ' and part number ' <i>Part Number</i> ' was installed in the system.
Troubleshooting	Was this PSU actually installed?

CPU_CHANGE_DETECTED

Severity	major
Description	<i>Component ID</i> has been changed from a serial of <i>old_serial</i> to <i>new_serial</i> .
Troubleshooting	Was this CPU actually replaced?

NIC_CHANGE_DETECTED

Severity	major
Description	<i>Component ID</i> has been changed from a serial of <i>old_serial</i> to <i>new_serial</i> .
Troubleshooting	Was this NIC actually replaced?

VPD_CHANGE_DETECTED

Severity	informational
Description	Change in VPD <i>VPD Name</i> to a value of ' <i>VPD Value</i> '.
Troubleshooting	NA. This information is for the event center.

MFG_CHANGE_DETECTED

Severity	informational
Description	Change in MFG <i>MFG Name</i> to a value of ' <i>MFG Value</i> '.
Troubleshooting	NA. This information is for the event center.

MM_CONFIG_CHANGE_DETECTED

Severity	informational
Description	Change in MM <i>MM Name</i> to a value of ' <i>MM Value</i> '.
Troubleshooting	NA. This information is for the event center.

TECHNICIAN_WORK_STARTED

Severity	informational
Description	Technician work has started, expected to end at <i>End Time</i> . Comment: <i>Comment</i> .
Troubleshooting	N/A

TECHNICIAN_WORK_ENDED

Severity	informational
Description	Technician work has ended after <i>Elapsed Time</i> minutes. Comment: <i>Comment</i> .
Troubleshooting	N/A

TECHNICIAN_WORK_TIMED_OUT

Severity	warning
Description	Technician work has timed out after <i>Elapsed Time</i> minutes. Comment: <i>Comment</i> .
Troubleshooting	N/A

XIV_SUPPORT_ENABLED

Severity	informational
Description	XIV support access from <i>From</i> is enabled from <i>Start Time</i> until <i>Finish Time</i> . Comment: <i>Comment</i> .
Troubleshooting	N/A

XIV_SUPPORT_ENABLED_NO_TIME_LIMIT

Severity	informational
Description	XIV support access from <i>From</i> is enabled from <i>Start Time</i> until explicitly disabled. Comment: <i>Comment</i> .

XIV_SUPPORT_DISABLED

Severity	informational
Description	XIV support access is disabled.
Troubleshooting	N/A

XIV_SUPPORT_WINDOW_TIMEOUT

Severity	informational
Description	XIV support work window timeout is expired.
Troubleshooting	N/A

HOST_DEFINE

Severity	informational
Description	Host of type <i>host.type</i> was defined with name ' <i>host.name</i> '.

HOST_UPDATE

Severity	informational
Description	Host named ' <i>host.name</i> ' was updated.

CLUSTER_CREATE

Severity	informational
Description	Cluster was defined with name ' <i>cluster.name</i> '.

HOST_DEFINE_FAILED_TOO_MANY

Severity	warning
Description	Host with name ' <i>name</i> ' could not be defined. You are attempting to define more hosts than the system permits.
Troubleshooting	Delete Hosts to allow new ones to be defined.

CLUSTER_CREATE_FAILED_TOO_MANY

Severity	warning
Description	Cluster with name ' <i>name</i> ' could not be defined. You are attempting to define more Clusters than the system permits.
Troubleshooting	Delete Clusters to allow new ones to be defined.

HOST_RENAME

Severity	informational
Description	Host with name ' <i>old_name</i> ' was renamed ' <i>host.name</i> '.

CLUSTER_RENAME

Severity	informational
Description	Cluster with name ' <i>old_name</i> ' was renamed ' <i>cluster.name</i> '.

HOST_DELETE

Severity	informational
Description	Host with name ' <i>host.name</i> ' was deleted.

CLUSTER_DELETE

Severity	informational
Description	Cluster with name ' <i>cluster.name</i> ' was deleted.

HOST_ADD_PORT

Severity	informational

Description	Port of type <i>type</i> and ID ' <i>port_name</i> ' was added to Host with name ' <i>host.name</i> '.

CLUSTER_ADD_HOST

Severity	informational
Description	Host with name ' <i>host.name</i> ' was added to Cluster with name ' <i>cluster.name</i> '.

HOST_REMOVE_PORT

Severity	informational
Description	Port of type <i>type</i> and ID ' <i>port_name</i> ' was removed from Host with name ' <i>host.name</i> ' was deleted.

CLUSTER_REMOVE_HOST

Severity	informational
Description	Host with name ' <i>host.name</i> ' was removed from Cluster with name ' <i>cluster.name</i> '.

DESTINATION_DEFINE

Severity	informational
Description	Destination with name ' <i>name</i> ' was defined.

DESTINATION_UPDATE

Severity	informational
Description	Destination with name ' <i>name</i> ' was updated.

DESTINATION_DELETE

Severity	informational
Description	Destination with name ' <i>name</i> ' was deleted.

DESTINATION_RENAME

Severity	informational
Description	Destination with name ' <i>old name</i> ' was renamed ' <i>new name</i> '.

DESTINATION_GROUP_CREATE

Severity	informational
Description	Destination Group with name ' <i>name</i> ' was created.

DESTINATION_GROUP_UPDATE

Severity	informational
Description	Destination Group with name ' <i>name</i> ' was updated.

DESTINATION_GROUP_DELETE

Severity	informational
Description	Destination Group with name ' <i>name</i> ' was deleted.

DESTINATION_GROUP_RENAME

Severity	informational
Description	Destination Group with name ' <i>old name</i> ' was renamed ' <i>new name</i> '.

DESTINATION_GROUP_ADD_DESTINATION

Severity	informational
Description	Destination with name ' <i>destination name</i> ' was added to destination group ' <i>destgroup name</i> '.

DESTINATION_GROUP_REMOVE_DESTINATION

Severity	informational
Description	Destination with name ' <i>destination name</i> ' was removed from destination group ' <i>destgroup name</i> '.

RULE_CREATE

Severity	informational
Description	Rule with name ' <i>name</i> ' was created.

RULE_UPDATE

Severity	informational
Description	Rule with name ' <i>name</i> ' was updated.

RULE_DELETE

Severity	informational
Description	Rule with name ' <i>name</i> ' was deleted.

RULE_RENAME

Severity	informational
Description	Rule with name ' <i>old name</i> ' was renamed ' <i>new name</i> '.

SMTP_GATEWAY_DEFINE

Severity	informational
Description	SMTP gateway with name ' <i>name</i> ' was defined.

SMTP_GATEWAY_UPDATE

Severity	informational
Description	SMTP gateway with name ' <i>name</i> ' was updated.

SMTP_GATEWAY_DELETE

Severity	informational
Description	SMTP gateway with name ' <i>name</i> ' was deleted.

SMTP_GATEWAY_RENAME

Severity	informational
Description	SMTP gateway with name ' <i>old name</i> ' was renamed ' <i>new name</i> '.

SMTP_GATEWAY_PRIORITIZE

Severity	informational
Description	SMTP gateways were prioritized; the new order is <i>order</i> .

SMTP_GATEWAY_FAILED

Severity	major
Description	SMTP gateway with name ' <i>name</i> ' has failed. It will not be used until <i>Retry Time</i> .

SMTP_GATEWAY_VIA_NODE_FAILED

Severity	warning
Description	Sending event <i>Event Code (Event Index)</i> through <i>SMTP Gateway</i> via <i>Module ID</i> has failed; Error message: ' <i>Error Message</i> '.

SMS_GATEWAY_DEFINE

Severity	informational
Description	SMS gateway with name ' <i>name</i> ' was defined.

SMS_GATEWAY_UPDATE

Severity	informational
Description	SMS gateway with name ' <i>name</i> ' was updated.

SMS_GATEWAY_DELETE

Severity	informational
Description	SMS gateway with name ' <i>name</i> ' was deleted.

SMS_GATEWAY_RENAME

Severity	informational
Description	SMS gateway with name ' <i>old name</i> ' was renamed ' <i>new name</i> '.

SMS_GATEWAY_PRIORITIZE

Severity	informational
Description	SMS gateways were prioritized; the new order is <i>order</i> .

EVENTS_WERE_DISCARDED

Severity	variable
Description	<i>Number of events</i> pending events of maximal severity ' <i>maximal severity</i> ' were discarded because of overload.
Troubleshooting	Please contact support.

CONS_GROUP_CREATE

Severity	informational
Description	Consistency Group with name ' <i>cg.name</i> ' was created.

CONS_GROUP_CREATE_FAILED_TOO_MANY

Severity	warning
Description	Consistency Group with name ' <i>name</i> ' could not be created. You are attempting to add more Consistency Groups than the system permits.
Troubleshooting	Delete Consistency Groups to allow new ones to be created.

CONS_GROUP_RENAME

Severity	informational
Description	Consistency Group with name ' <i>old_name</i> ' was renamed ' <i>cg.name</i> '.

CONS_GROUP_DELETE

Severity	informational
Description	Consistency Group with name ' <i>cg.name</i> ' was deleted.

CONS_GROUP_ADD_VOLUME

Severity	informational
Description	Volume with name ' <i>volume.name</i> ' was added to Consistency Group with name ' <i>cg.name</i> '.

SLAVE_CONS_GROUP_ADD_VOLUME

Severity	informational
Description	Volume with name ' <i>volume.name</i> ' was added to Consistency Group with name ' <i>cg.name</i> ' by its remote peer.

CONS_GROUP_REMOVE_VOLUME

Severity	informational
Description	Volume with name ' <i>volume.name</i> ' was removed from Consistency Group with name ' <i>cg.name</i> '.

SLAVE_CONS_GROUP_REMOVE_VOLUME

Severity	informational
Description	Volume with name ' <i>volume.name</i> ' was removed from Consistency Group with name ' <i>cg.name</i> ' by its remote peer.

CONS_GROUP_SNAPSHOTS_CREATE

Severity	informational
Description	Snapshot Group for Consistency Group with name ' <i>cg.name</i> ' was created with name ' <i>cs_name</i> '.

CONS_GROUP_SNAPSHOTS_CREATE_FAILED_TOO_MANY

Severity	warning
Description	Snapshot Group for Consistency Group ' <i>cg.name</i> ' could not be created. You are attempting to add more Volumes than the system permits.
Troubleshooting	Delete Volumes to allow new ones to be created.

CONS_GROUP_SNAPSHOTS_OVERWRITE

Severity	informational
Description	Snapshot Group named ' <i>cs_name</i> ' was overridden for Consistency Group with name ' <i>cg.name</i> '.

SLAVE_CONS_GROUP_SNAPSHOTS_CREATE

Severity	informational
Description	Mirrored Snapshot Group for Consistency Group with name ' <i>cg.name</i> ' was created with name ' <i>cs_name</i> '.

SLAVE_CONS_GROUP_SNAPSHOTS_OVERWRITE

Severity	informational
Description	Mirrored Snapshot Group named ' <i>cs_name</i> ' was overridden for Consistency Group with name ' <i>cg.name</i> '.

MIRROR_CONS_GROUP_SNAPSHOTS_CREATE

Severity	informational
Description	Mirrored Snapshot Group for Consistency Group with name ' <i>cg.name</i> ' was created with name ' <i>cs_name</i> '.

MIRROR_CONS_GROUP_SNAPSHOTS_OVERWRITE

Severity	informational
Description	Mirrored Snapshot Group named ' <i>cs_name</i> ' was overridden for Consistency Group with name ' <i>cg.name</i> '.

MIRROR_SNAPGROUP_CREATE_FAILED

Severity	minor
Description	Remote snapshot group named ' <i>snapshot group name</i> ' was not created successfully. Error code is ' <i>error</i> '.

SNAPSHOT_GROUP_RESTORE

Severity	informational
Description	Volumes were restored from Snapshot Group with name ' <i>cs_name</i> '.

SNAPSHOT_GROUP_RENAME

Severity	informational
Description	Snapshot Group with name ' <i>cs_name</i> ' were renamed to ' <i>new_name</i> '.

SNAPSHOT_GROUP_DUPLICATE

Severity	informational
Description	All Snapshots in Snapshot Group with name ' <i>cs_name</i> ' were duplicated. Duplicate Snapshot Group is named ' <i>new_cs_name</i> '.

SNAPSHOT_GROUP_FORMAT

Severity	informational
Description	All Snapshots in Snapshot Group with name ' <i>cs_name</i> ' were formatted.

SNAPSHOT_GROUP_DELETE

Severity	informational
Description	All Snapshots in Snapshot Group with name ' <i>cs_name</i> ' were deleted.

SNAPSHOT_GROUP_CHANGE_PRIORITY

Severity	informational
Description	Deletion Priority of all Snapshots in Snapshot Group with name ' <i>cs_name</i> ' were changed from ' <i>old priority</i> ' to ' <i>new priority</i> '.

SNAPSHOT_GROUP_LOCK

Severity	informational
Description	All Snapshots in Snapshot Group with name ' <i>cs_name</i> ' were locked.

SNAPSHOT_GROUP_UNLOCK

Severity	informational
Description	All Snapshots in Snapshot Group with name ' <i>cs_name</i> ' were unlocked.

SNAPSHOT_GROUP_DELETED_DUE_TO_POOL_EXHAUSTION

Severity	informational
Description	All Snapshots in Snapshot Group with name ' <i>snapshot.sg_name</i> ' have been deleted because Storage Pool with name ' <i>snapshot.pool_name</i> ' is full.
Troubleshooting	N/A

SNAPSHOT_GROUP_DISBAND

Severity	informational
Description	Snapshot Group with name ' <i>cs_name</i> ' was dismantled. All Snapshots which belonged to that Snapshot Group should be accessed directly.

CONS_GROUP_MOVE

Severity	informational
Description	Consistency Group with name ' <i>cg.name</i> ' has been moved from Storage Pool ' <i>orig_pool.name</i> ' to Pool ' <i>pool.name</i> '.

CONS_GROUP_GROUPED_POOL_MOVE

Severity	informational
Description	Consistency Group with name ' <i>cg.name</i> ' has been moved from Grouped Pool ' <i>orig_gp.name</i> ' to Grouped Pool ' <i>gp.name</i> '.

XCG_CREATE

Severity	informational
Description	Cross Consistency Group with name ' <i>xcg</i> ' was created.

XCG_DELETE

Severity	informational
Description	Cross Consistency Group with name ' <i>xcg</i> ' was deleted.

XCG_ADD_CG

Severity	informational
Description	CG with name ' <i>cg.name</i> ' was added to Cross Consistency Group with name ' <i>xcg</i> '.

XCG_REMOVE_CG

Severity	informational
Description	CG with name ' <i>cg.name</i> ' was removed from Cross Consistency Group with name ' <i>xcg</i> '.

GROUPED_POOL_CREATE

Severity	informational
Description	Grouped Pool with name ' <i>gp.name</i> ' was created.

GROUPED_POOL_RENAME

Severity	informational
Description	Grouped Pool with name ' <i>old_name</i> ' was renamed ' <i>gp.name</i> '.

GROUPED_POOL_DELETE

Severity	informational
Description	Grouped Pool with name ' <i>gp.name</i> ' was deleted.

GROUPED_POOL_CAPACITY_SHIFT

Severity	informational
Description	On Grouped Pool with name ' <i>gp.name</i> ' Capacity of <i>capacity_size</i> GB was shifted from pool ' <i>src_pool.name</i> ' to pool ' <i>dest_pool.name</i> '.

TARGET_DEFINE

Severity	informational
Description	Target was defined named ' <i>target.name</i> '.

TARGET_DEFINE_FAILED_TOO_MANY

Severity	warning
Description	Target could not be defined. You are attempting to define more targets than the system permits.
Troubleshooting	Delete targets to allow new ones to be defined.

TARGET_RENAME

Severity	informational
Description	Target named ' <i>old_name</i> ' was renamed ' <i>target.name</i> '.

TARGET_DELETE

Severity	informational
Description	Target named ' <i>target.name</i> ' was deleted.

TARGET_ALLOW_ACCESS

Severity	informational
Description	Target ' <i>target.name</i> ' is allowed to access this machine.

TARGET_PORT_ADD

Severity	informational
Description	Port ' <i>port_name</i> ' was added to target named ' <i>target.name</i> '.

TARGET_PORT_REMOVE

Severity	informational
Description	Port ' <i>port_name</i> ' was removed from target named ' <i>target.name</i> '.

TARGET_PORT_ACTIVATE

Severity	informational
Description	Port ' <i>port_name</i> ' in target named ' <i>target.name</i> ' was activated.

TARGET_PORT_DEACTIVATE

Severity	informational
Description	Port ' <i>port_name</i> ' was deactivated in target named ' <i>target.name</i> '.

TARGET_CONNECTIVITY_CREATE

Severity	informational
Description	Port ' <i>Connection Remote Port Address</i> ' of target named ' <i>Connection Target Name</i> ' is connected to the system through <i>Local FC Port</i> .

TARGET_ISCSI_CONNECTIVITY_CREATE

Severity	informational
Description	Port ' <i>Connection Remote Port Address</i> ' of target named ' <i>Connection Target Name</i> ' is connected to the system through ip interface ' <i>Local IP interface</i> '.

TARGET_CONNECTIVITY_CREATE_FAILED_TOO_MANY

Severity	warning
Description	Port could not be connected to the system. You are attempting to define more connections than the system permits.
Troubleshooting	Delete Connections to allow new ones to be created.

TARGET_CONNECTIVITY_DELETE

Severity	informational
Description	Port ' <i>Connection Remote Port Address</i> ' of target named ' <i>Connection Target Name</i> ' was disconnected from <i>Local FC Port</i> .

TARGET_ISCSI_CONNECTIVITY_DELETE

Severity	informational
Description	Port ' <i>Connection Remote Port Address</i> ' of target named ' <i>Connection Target Name</i> ' was disconnected from ip interface ' <i>Local IP interface</i> '.

TARGET_CONNECTIVITY_ACTIVATE

Severity	informational
Description	Connectivity between Port ' <i>Connection Remote Port Address</i> ' of target named ' <i>Connection Target Name</i> ' and <i>Local FC Port</i> was activated.

TARGET_ISCSI_CONNECTIVITY_ACTIVATE

Severity	informational
Description	Connectivity between Port ' <i>Connection Remote Port Address</i> ' of target named ' <i>Connection Target Name</i> ' and ip interface ' <i>Local IP interface</i> ' was activated.

TARGET_CONNECTIVITY_DEACTIVATE

Severity	informational
Description	Connectivity between Port ' <i>Connection Remote Port Address</i> ' of target named ' <i>Connection Target Name</i> ' and <i>Local FC Port</i> was deactivated.

TARGET_ISCSI_CONNECTIVITY_DEACTIVATE

Severity	informational
Description	Connectivity between Port ' <i>Connection Remote Port Address</i> ' of target named ' <i>Connection Target Name</i> ' and ip interface ' <i>Local IP interface</i> ' was deactivated.

TARGET_CONNECTIVITY_CONFLICT_DETECTED

Severity	major
Description	Connectivity between Port ' <i>Connection Remote Port Address</i> ' of target named ' <i>Connection Target Name</i> ' and FC port ' <i>Local IP interface</i> ' will be deleted due to a connectivity conflict.

TARGET_ISCSI_CONNECTIVITY_CONFLICT_DETECTED

Severity	major
Description	Connectivity between Port ' <i>Connection Remote Port Address</i> ' of target named ' <i>Connection Target Name</i> ' and IP interface ' <i>Local IP interface</i> ' will be deleted due to a connectivity conflict.

TARGET_CONNECTION_ESTABLISHED

Severity	informational
Description	Target named ' <i>target.name</i> ' is accessible through remote service <i>module_id</i> .

TARGET_CONNECTION_DISCONNECTED

Severity	minor
Description	Target named ' <i>target.name</i> ' is no longer accessible through remote service <i>module_id</i> .

TARGET_DISCONNECTED

Severity	major
Description	Target named ' <i>target.name</i> ' is no longer accessible through any gateway module.

TARGET_CLOCK_SKEW_ABOVE_LIMIT

Severity	warning
Description	Target ' <i>target.name</i> ' has clock skew above the allowed limit relative to local machine.

TARGET_CLOCK_SKEW_RESOLVED

Severity	informational
Description	Target named ' <i>target.name</i> ' clock skew has been resolved.

TARGET_LINK_DOWN_BEYOND_THRESHOLD

Severity	major
Description	Target named ' <i>target.name</i> ' is not accessible for a long time.

OLVM_DELETE_ALL_REFERENCES_TO_SOURCE

Severity	major
Description	Target named ' <i>target.name</i> ' was released from all IBM Hyper-Scale Mobility relationships.

SNAPSHOT_CREATE

Severity	informational
Description	Snapshot named ' <i>snapshot.name</i> ' was created for volume named ' <i>volume.name</i> '.

SNAPSHOT_CREATE_MANY

Severity	informational
Description	Created <i>num_of_vols</i> snapshots.

SNAPSHOT_OVERWRITE

Severity	informational
Description	Snapshot named ' <i>snapshot.name</i> ' was overridden for volume named ' <i>volume.name</i> '.

SNAPSHOT_FORMAT

Severity	informational
Description	Snapshot named ' <i>snapshot.name</i> ' was formatted.

SNAPSHOT_CREATE_FAILED_TOO_MANY

Severity	warning
Description	Snapshot for volume named ' <i>volume.name</i> ' could not be created. You are attempting to add more volumes than the system permits.
Troubleshooting	Delete volumes to allow new ones to be created.

SNAPSHOT_DUPLICATE

Severity	informational
Description	Snapshot named ' <i>snapshot.name</i> ' was created as duplicate of Snapshot named ' <i>original_snapshot.name</i> '.

SNAPSHOT_DUPLICATE_FAILED_TOO_MANY

Severity	warning
Description	Snapshot named ' <i>snapshot.name</i> ' could not be duplicated. You are attempting to add more volumes than the system permits.
Troubleshooting	Delete volumes to allow new ones to be created.

SNAPSHOT_RESTORE

Severity	informational
Description	Volume named ' <i>volume.name</i> ' was restored from Snapshot named ' <i>snapshot.name</i> '.

SNAPSHOT_CHANGE_PRIORITY

Severity	informational
Description	Snapshot Delete Priority of Snapshot named ' <i>snapshot.name</i> ' was changed from ' <i>old_priority</i> ' to ' <i>snapshot.delete_priority</i> '.
Troubleshooting	N/A

SNAPSHOT_DELETED_DUE_TO_POOL_EXHAUSTION

Severity	warning
Description	Snapshot named ' <i>snap.name</i> ' has been deleted because Storage Pool named ' <i>snap.pool_name</i> ' is full.
Troubleshooting	N/A

MIRROR_SNAPSHOT_CREATE

Severity	informational
Description	Mirrored Snapshot named ' <i>snapshot.name</i> ' was created for volume named ' <i>volume.name</i> '.

MIRROR_SNAPSHOT_CREATE_FAILED

Severity	minor
Description	Remote snapshot named ' <i>snapshot name</i> ' was not created successfully. Error code is ' <i>error</i> '.

MIRROR_SNAPSHOT_OVERWRITE

Severity	informational
Description	Mirrored Snapshot named ' <i>snapshot.name</i> ' was overridden for volume named ' <i>volume.name</i> '.

MIRROR_SLAVE_SNAPSHOT_CREATE

Severity	informational
Description	Mirrored Snapshot named ' <i>snapshot.name</i> ' was created for volume named ' <i>volume.name</i> '.

MIRROR_SLAVE_SNAPSHOT_OVERWRITE

Severity	informational
Description	Mirrored Snapshot named ' <i>snapshot.name</i> ' was overridden for volume named ' <i>volume.name</i> '.

INVALID_ASYNC_ASSOC

Severity	critical
Description	Can not start async job without next job. volume uid <i>volume uid</i>
Troubleshooting	N/A

INVALID_DATA_GENERATION_ID

Severity	minor
Description	Rejected write with id <i>write id</i> on volume <i>volume nr</i> for lba <i>lba</i> . Expected id is <i>expected id</i>
Troubleshooting	N/A

UNUSUAL_CONF_LOCK_TIME

Severity	warning
Description	very long conf-lock duration. node= <i>node</i> , duration_msec= <i>duration</i> ,
Troubleshooting	look at traces for more details.

XDRP_FLAGS_ON_AND_NO_MIRROR

Severity	critical
Description	partition <i>partition number</i> on volume <i>Volume</i> has xdrp flags but master volume <i>master volume</i> has no mirroring.
Troubleshooting	N/A

FAILED_TO_ALLOC_FOR_REMOTE_FLAG

Severity	critical
Description	Failed to allocate partition number <i>partition number</i> on volume <i>volume</i> disk <i>disk id</i> when trying to mark bits for remote sync.
Troubleshooting	N/A

FAILED_ALLOC_IN_REBUILD

Severity	critical
Description	Failed to allocate partition number <i>Partition Number</i> on volume <i>volume number</i> disk <i>disk id</i> during REBUILD.
Troubleshooting	N/A

ACQUIRED_INVALID_PARTITION

Severity	critical
Description	Acquired invalid partition number <i>Partition Number</i> on volume <i>volume number</i> , flags <i>partition flags</i> reference count <i>reference count</i> .
Troubleshooting	N/A

ACQUIRED_INVALID_PARTIION

Severity	critical
Description	Acquired invalid partition number <i>Partition Number</i> on volume <i>volume number</i> , flags <i>partition flags</i> reference count <i>reference count</i> .
Troubleshooting	N/A

INVALID_PARTITION_STATIC_FLAG

Severity	critical
Description	Partition static flag doesn't match requested metadata static flag. Partition number <i>Partition number</i> volume <i>volume number</i> flags <i>partition flags</i> metadata flags <i>metadata flags</i> .
Troubleshooting	N/A

SUSPECT_DATA_LOSS

Severity	critical
Description	Suspected data loss on Partition <i>Disk ID</i> , volume= <i>Volume</i> , logical-partition= <i>Logical Partition Number</i> , physical-partition= <i>Physical Partition Number</i> .
Troubleshooting	Verify that data is unreadable. Use backup data to recover.

SCRUBBING_CHECKSUM_DIFF

Severity	critical
Description	Scrubbing checksum diff. Primary: <i>Primary Disk</i> , checksum= <i>Primary Checksum</i> p_phy_part_nr= <i>p_phy_part_nr</i> ; Secondary: <i>Secondary Disk</i> checksum= <i>Secondary Checksum</i> s_phy_part_nr= <i>s_phy_part_nr</i> ; Volume= <i>Volume</i> , partition= <i>Logical Partition Number</i> , is_master= <i>Is Master</i> .
Troubleshooting	Compare data on primary and secondary.

SCRUBBER_FIXED_CHECKSUM_DIFF

Severity	informational
Description	Scrubber fixed checksum diff. Volume= <i>Volume</i> , Partition= <i>Logical Partition Number</i> . Fixed from node= <i>Node</i> , disk= <i>Disk</i>
Troubleshooting	N/A

SCRUBBING_CHECKSUM_DIFF_RETRY_COUNT

Severity	critical
Description	Scrub passed partition <i>number of retries</i> times and found diffs <i>number of times diff found</i> times. primary: <i>Primary Disk</i> secondary: <i>Secondary Disk</i> volume= <i>Volume</i> , partition= <i>Logical Partition Number</i> .
Troubleshooting	Compare data on primary and secondary.

SCRUBBING_REMOTE_DIGEST_DIFF

Severity	critical
Description	Scrubbing found different digests in local and remote.disk <i>disk Tracks Diff Count</i> tracks are different. First diff track <i>Track</i> local: (<i>Local Digest 0</i> , <i>Local Digest 1</i>) Remote: (<i>Remote Digest 0</i> , <i>Remote Digest 1</i>), volume= <i>Volume</i> , pvi= <i>PVI</i> , partition= <i>Logical Partition Number</i> .
Troubleshooting	Compare data on primary and secondary.

SCRUBBING_REMOTE_DIGEST_DIFF_IDENTIFIED_AND_CORRECTED

Severity	informational
Description	Scrubbing identified and corrected data on remote system. disk <i>disk Tracks Diff Count</i> tracks were different. First diff track <i>Track</i> local: (<i>Local Digest 0</i> , <i>Local Digest 1</i>) Remote: (<i>Remote Digest 0</i> , <i>Remote Digest 1</i>), volume= <i>Volume</i> , pvi= <i>PVI</i> , partition= <i>Logical Partition Number</i> .
Troubleshooting	N/A

SCRUBBING_REMOTE_DIGEST_DIFF_CORRECT_FAILED

Severity	informational
Description	Scrubbing failed to correct remote system. disk <i>disk Tracks Diff Count</i> tracks are different. First diff track <i>Track</i> local: (<i>Local Digest 0</i> , <i>Local Digest 1</i>) Remote: (<i>Remote Digest 0</i> , <i>Remote Digest 1</i>), volume= <i>Volume</i> , pvi= <i>PVI</i> , partition= <i>Logical Partition Number</i> .
Troubleshooting	N/A

SCRUBBING_FORMATTED_NOT_ZERO

Severity	critical
Description	Scrubbing found formatted partition with non zero checksum on <i>Disk ID</i> , partition= <i>phy_part_nr</i> , checksum= <i>checksum</i> .
Troubleshooting	N/A

SCRUBBING_SUCCESS

Severity	informational
Description	Scrubbing found that digests are now equal at location: <i>Primary Disk</i> , p_phy_part_nr= <i>p_phy_part_nr</i> ; Secondary: <i>Secondary Disk</i> s_phy_part_nr= <i>s_phy_part_nr</i> ; Volume= <i>Volume</i> , partition= <i>Logical Partition Number</i> .
Troubleshooting	N/A

SCRUBBING_SECONDARY_DIFF_IDENTIFIED_AND_CORRECTED

Severity	informational
Description	Scrubbing identified and corrected diff on secondary: <i>Secondary Disk</i> s_phy_part_nr= <i>s_phy_part_nr</i> ; From primary: <i>Primary Disk</i> , p_phy_part_nr= <i>p_phy_part_nr</i> ; Volume= <i>Volume</i> , partition= <i>Logical Partition Number</i> .
Troubleshooting	N/A

MEDIUM_ERROR_IN_DATA_MIGRATION

Severity	critical
Description	Medium error in data migration into volume ' <i>Volume Name</i> ' at LBA <i>LBA</i> for <i>Length</i> blocks.
Troubleshooting	Remote machine indicated Medium Error when read.

ZERO_LENGTH_IO

Severity	warning
Description	Media errors on node= <i>node</i> , interface= <i>interface</i> , volume= <i>volume</i> , LBA= <i>LBA</i> , blk_cnt= <i>Block Count</i> .
Troubleshooting	N/A

BUFFER_POOL_EMPTY

Severity	major
Description	The memory pool of cache buffers in node= <i>node</i> is exhausted. All <i>pool_size</i> buffers are allocated.
Troubleshooting	N/A

ORPHANED_ASYNC_PARTITION

Severity	warning
Description	An active async job (<i>Job ID</i>) running on cache <i>node</i> disk <i>Disk ID</i> found partition <i>Logical partition number</i> of source and target snapshots to be the same partition. Master volume= <i>Master Volume</i> .
Troubleshooting	N/A

CACHE_ALLOCATION_NO_READ_CACHE

Severity	critical
Description	Cache had no read-cache pages available for allocation in the context of <i>disk</i> .
Troubleshooting	Look at traces for more details.

CACHE_COMPRESSION_MEMORY_ALLOCATION_FAILURE

Severity	critical
Description	Cache had no read-cache pages available for compression memory allocation in the context of <i>Disk ID</i> . Required= <i>Required</i> , available= <i>available</i> .
Troubleshooting	Look at traces for more details.

SSD_GET_KEY_FAILED

Severity	warning
Description	Failed to get key for SSD <i>SSD ID</i>
Troubleshooting	N/A

SSD_GET_KEY_COMPLETED

Severity	informational
Description	Get key completed for SSD <i>SSD ID</i>
Troubleshooting	N/A

SSD_CACHE_UPDATED_TO_NEW_MD_VERSION

Severity	informational
Description	SSD cache of <i>SSD ID</i> was reset and updated to new metadata version
Troubleshooting	N/A

SSD_DATA_INTEGRITY_ERROR_DETECTED

Severity	minor
Description	Read from SSD <i>Disk ID</i> failed the integrity check due to <i>Reason</i> , Page Number= <i>Page Number</i> . Read from disk instead.
Troubleshooting	N/A

SSD_ERROR_DETECTED

Severity	minor
Description	Read from SSD <i>Disk ID</i> failed the integrity check due to <i>Reason</i> , Page Number= <i>Page Number</i> . Read from disk instead.
Troubleshooting	N/A

DISK_SHOULD_FAIL

Severity	major
Description	<i>Disk ID</i> is malfunctioning and should fail.
Troubleshooting	Please contact your Administrator.

DISK_NEEDS_PHASEOUT

Severity	major
Description	<i>Disk ID</i> needs to be phased out.
Troubleshooting	Please contact your Administrator.

SSD_DISK_LABELS_MISMATCH

Severity	major
Description	SSD <i>SSD ID</i> has data that mismatches disk <i>Disk ID</i>

IGNORED_ALLOCATED_NOT_USED_PARTITION

Severity	informational
Description	While looking for partition (<i>Disk</i> ., <i>Vol</i> ., <i>Part</i> .) we unexpectedly encountered a partition higher in tree (<i>Disk</i> ., <i>Vol</i> ., <i>Part</i> .).

SCRUBBER_WAS_REPOSITIONED

Severity	major
Description	The position of the scrubber on disk <i>Disk ID</i> was modified by the user (from <i>Old Value</i> to <i>New Value</i>).
Troubleshooting	

STATUS_AGENT_ERROR

Severity	critical
Description	Status Agent error: <i>message</i> (value= <i>value</i>).
Troubleshooting	Please contact support.

USER_DEFINED

Severity	informational
Description	A user with name ' <i>Name</i> ' and category <i>Category</i> was defined.
Troubleshooting	N/A

USER_DELETED

Severity	informational
Description	A user with name ' <i>Name</i> ' and category <i>Category</i> was deleted.
Troubleshooting	N/A

USER_RENAMED

Severity	informational
Description	User with name ' <i>Old Name</i> ' was renamed ' <i>New Name</i> '.
Troubleshooting	N/A

USER_UPDATED

Severity	informational
Description	User with name ' <i>Name</i> ' was updated.
Troubleshooting	N/A

USER_ADDED_TO_USER_GROUP

Severity	informational
Description	User ' <i>User Name</i> ' was added to user group ' <i>User Group Name</i> '.
Troubleshooting	N/A

USER_REMOVED_FROM_USER_GROUP

Severity	informational
Description	User ' <i>User Name</i> ' was removed from user group ' <i>User Group Name</i> '.
Troubleshooting	N/A

USER_GROUP_CREATED

Severity	informational
Description	A user group with name ' <i>Name</i> ' was created.
Troubleshooting	N/A

USER_GROUP_DELETED

Severity	informational
Description	A user group with name ' <i>Name</i> ' was deleted.
Troubleshooting	N/A

USER_GROUP_RENAMED

Severity	informational
Description	User group with name ' <i>Old Name</i> ' was renamed ' <i>New Name</i> '.
Troubleshooting	N/A

LDAP_AUTHENTICATION_ACTIVATED

Severity	informational
Description	LDAP authentication activated.
Troubleshooting	N/A

LDAP_AUTHENTICATION_DEACTIVATED

Severity	warning
Description	LDAP authentication deactivated.
Troubleshooting	N/A

LDAP_CONFIGURATION_CHANGED

Severity	warning
Description	LDAP configuration has changed.
Troubleshooting	N/A

LDAP_CONFIGURATION_RESET

Severity	warning
Description	LDAP configuration has reset.
Troubleshooting	N/A

USER_LOGIN_HAS_SUCCEEDED

Severity	informational
Description	User ' <i>User Name</i> ' from IP ' <i>Client Address</i> ' successfully logged into the system.
Troubleshooting	N/A

USER_LOGIN_HAS_FAILED

Severity	warning
Description	User ' <i>User Name</i> ' from IP ' <i>Client Address</i> ' failed logging into the system.
Troubleshooting	N/A

USER_HAS_FAILED_TO_RUN_COMMAND

Severity	warning
Description	User ' <i>User Name</i> ' from IP ' <i>Client Address</i> ' failed authentication when trying to run command ' <i>Command Line</i> '.
Troubleshooting	N/A

LDAP_SERVER_INACCESSIBLE

Severity	minor
Description	LDAP server <i>FQDN</i> is inaccessible.
Troubleshooting	N/A

LDAP_SERVER_ACCESSIBLE

Severity	informational
Description	LDAP server <i>FQDN</i> is now accessible.
Troubleshooting	N/A

LDAP_SSL_CERTIFICATE_ABOUT_TO_EXPIRE

Severity	warning
Description	SSL Certificate of LDAP server ' <i>Server FQDN</i> ' is about to expire on <i>Expiration Date</i> (Counter notification).
Troubleshooting	N/A

LDAP_SERVER_WAS_ADDED

Severity	informational
Description	LDAP server ' <i>Server FQDN</i> ' was added to the system.

Troubleshooting	N/A

LDAP_SERVER_WAS_REMOVED

Severity	informational
Description	LDAP server ' <i>Server FQDN</i> ' was removed from the system.
Troubleshooting	N/A

DESIGNATED_MSM_USER

Severity	informational
Description	<i>Description</i>
Troubleshooting	N/A

DOMAIN_POLICY_SET

Severity	informational
Description	Domain policy for <i>Parameter Name</i> set to ' <i>Parameter Value</i> '
Troubleshooting	N/A

USER_ADDED_TO_DOMAIN

Severity	informational
Description	User <i>User Name</i> was added to domain <i>Domain Name (Exclusive)</i> .
Troubleshooting	N/A

USER_REMOVED_FROM_DOMAIN

Severity	informational
Description	User <i>User Name</i> was removed from domain <i>Domain Name</i> .
Troubleshooting	N/A

APPADMIN_CAPABILITIES_SET

Severity	informational
Description	Application admin capabilities have been set to <i>Capabilities</i>
Troubleshooting	N/A

ACCESS_TO_HOST_GRANTED_TO_USER_GROUP

Severity	informational
Description	User group ' <i>User Group Name</i> ' was granted access to host ' <i>Host Name</i> '.
Troubleshooting	N/A

ACCESS_OF_USER_GROUP_TO_HOST_REMOVED

Severity	informational
Description	Access of User group ' <i>User Group Name</i> ' to host ' <i>Host Name</i> ' was removed.
Troubleshooting	N/A

ACCESS_TO_CLUSTER_GRANTED_TO_USER_GROUP

Severity	informational
Description	User group ' <i>User Group Name</i> ' was granted access to cluster ' <i>Cluster Name</i> '.
Troubleshooting	N/A

ACCESS_OF_USER_GROUP_TO_CLUSTER_REMOVED

Severity	informational
Description	Access of User group ' <i>User Group Name</i> ' to cluster ' <i>Cluster Name</i> ' was removed.
Troubleshooting	N/A

COMPONENT_TEST_HAS_FAILED

Severity	minor
Description	Test of <i>Component ID</i> has failed. Failure reason: <i>Failure Reason</i> .
Troubleshooting	Please contact support.

COMPONENT_TEST_SUCCEEDED

Severity	informational
Description	Test of <i>Component ID</i> succeeded.
Troubleshooting	Please contact support.

MODULE_COMPONENT_TEST_STARTED

Severity	informational
Description	Test of <i>Component ID</i> started.
Troubleshooting	N/A

DISK_COMPONENT_TEST_STARTED

Severity	informational
Description	Test of <i>Component ID</i> started.
Troubleshooting	N/A

IB_SWITCH_COMPONENT_TEST_STARTED

Severity	informational
Description	Test of <i>Component ID</i> started.
Troubleshooting	N/A

SSD_COMPONENT_TEST_STARTED

Severity	informational
Description	Test of <i>Component ID</i> started.
Troubleshooting	N/A

COMPONENT_WAS_PHASED_OUT

Severity	informational
Description	<i>Component ID</i> was phased-out.
Troubleshooting	N/A

COMPONENT_WAS_FAILED

Severity	informational
Description	Component <i>Component ID</i> was marked as failed.
Troubleshooting	N/A

COMPONENT_FAILURE_WAS_CANCELED

Severity	informational
Description	Component <i>Component ID</i> failure status was reset.
Troubleshooting	N/A

COMPONENT_WAS_PHASED_IN

Severity	informational
Description	<i>Component ID</i> was phased-in.
Troubleshooting	N/A

COMPONENT_WAS_EQUIPPED

Severity	informational
Description	<i>Component ID</i> was equipped.
Troubleshooting	N/A

COMPONENT_WAS_UNEQUIPPED

Severity	informational
Description	<i>Component ID</i> was unequipped.
Troubleshooting	N/A

INTERFACE_SERVICES_ACTIVATED

Severity	informational
Description	Interface services of <i>Module ID</i> were activated.
Troubleshooting	N/A

INTERFACE_SERVICES_DEACTIVATED

Severity	informational
Description	Interface services of <i>Module ID</i> were deactivated.
Troubleshooting	N/A

COMPONENT_FIRMWARE_UPGRADE_ABORTING

Severity	warning
Description	Aborting <i>Upgrade type</i> upgrade of <i>Firmware type</i> firmware, version <i>Label</i> , on <i>Scope</i> . Abort reason: <i>Reason</i> . Waiting for current upgrade item to complete.
Troubleshooting	N/A

COMPONENT_FIRMWARE_UPGRADE_ABORTED

Severity	warning
Description	Aborted <i>Upgrade type</i> upgrade of <i>Firmware type</i> firmware, version <i>Label</i> , on <i>Scope</i> . Abort reason: <i>Reason</i> . Progress <i>Attempted/Total</i> , <i>Successes</i> succeeded, <i>Failures</i> failed, <i>No-Ops</i> no-ops.
Troubleshooting	N/A

COMPONENT_FIRMWARE_UPGRADE_DONE

Severity	informational
Description	Finished <i>Upgrade type</i> upgrade of <i>Firmware type</i> firmware, version <i>Label</i> , on <i>Scope</i> . <i>Successes</i> succeeded, <i>Failures</i> failed, <i>No-Ops</i> no-ops.
Troubleshooting	N/A

COMPONENT_FIRMWARE_UPGRADE_STARTED

Severity	informational
Description	Starting <i>Upgrade type</i> upgrade of <i>Firmware type</i> firmware, version <i>Label</i> , on <i>Scope</i> .
Troubleshooting	N/A

COMPONENT_FIRMWARE_CANNOT_PHASEOUT_COMPONENT

Severity	minor
Description	Cannot phase out <i>Component ID</i> : <i>Error</i> . Firmware upgrade result was: <i>Upgrade result</i> .
Troubleshooting	N/A

COMPONENT_FIRMWARE_CANNOT_FAIL_COMPONENT

Severity	minor
Description	Cannot fail <i>Component ID</i> : <i>Error</i> . Firmware upgrade result was: <i>Upgrade result</i> .
Troubleshooting	N/A

MIRRORING_CONNECTIVITY_TO_NON_XIV_TARGET

Severity	warning
Description	Gateway Node <i>#Node ID</i> : connection to <i>target name:target's connection index</i> mirroring connection was established, but being ignored because the remote end is not an XIV target or is not properly configured
Troubleshooting	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 <i>target_port</i> defined for this system.

DM_CONNECTIVITY_TO_XIV_TARGET

Severity	warning
Description	Gateway Node # <i>Node ID</i> : connection to <i>target name:target's connection index</i> DM connection was established, but being ignored because the remote end is an XIV target configured for mirroring, rather than a host
Troubleshooting	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 <i>target_port</i>).

TAKEN_OVER

Severity	informational
Description	Module <i>Module ID</i> has taken over as <i>Singleton Node ID</i> .

EMERGENCY_ROOT_ACCESS

Severity	warning
Description	Emergency login to 'root' account on module ' <i>Component ID</i> ' from ' <i>IP Address</i> ' using key number ' <i>Authorized Key Number</i> '.
Troubleshooting	N/A

EMERGENCY_CONSOLE_ACCESS

Severity	warning
Description	Emergency login to ' <i>Unix Account Name</i> ' account on module ' <i>Component ID</i> ' from tty ' <i>TTY Device</i> '.
Troubleshooting	N/A

CR_BYPASS_ACCESS

Severity	warning
Description	<i>Command</i> that bypasses CR mechanism access to ' <i>Unix Account Name</i> ' account on module ' <i>Component ID</i> ' from ' <i>IP Address</i> '.
Troubleshooting	N/A

CR_KEY_SETUP_OK

Severity	informational
Description	Challenge-response key was successfully set on all modules in the system.
Troubleshooting	N/A

CR_KEY_UPGRADE_NOT_DONE

Severity	warning
Description	Challenge-response key was not upgraded on the system since a valid key has been previously set.
Troubleshooting	N/A

CR_KEY_SETUP_FAILED

Severity	major
Description	Failed to set challenge-response key on module ' <i>Component ID</i> '.
Troubleshooting	N/A

SSH_REVOKE_KEY_OK

Severity	informational
Description	Authorized SSH key ending with ' <i>Tail of Authorized SSH key</i> ' was successfully revoked for user ' <i>Unix Account Name</i> ' on all modules in the system.
Troubleshooting	N/A

SSH_REVOKE_KEY_FAILED

Severity	major
Description	Failed to revoke authorized SSH key ending with ' <i>Tail of Authorized SSH key</i> ' for user ' <i>Unix Account Name</i> ' on module ' <i>Component ID</i> '.
Troubleshooting	N/A

CLEAR_FACTORY_SETTINGS_STARTED

Severity	informational
Description	Started clear factory settings.
Troubleshooting	Please contact support.

CLEAR_FACTORY_SETTINGS_FINISHED

Severity	informational
Description	Finished clear factory settings.
Troubleshooting	Please contact support.

CLEAR_FACTORY_SETTINGS_FAILED

Severity	critical
Description	Failed to finish clear factory settings.
Troubleshooting	Please contact support.

PROBLEMATIC_DISK_BEHAVIOR_DETECTED

Severity	variable
Description	<i>Component ID</i> exhibits problematic behavior.
Troubleshooting	Please contact support.

PROBLEMATIC_DISK_BEHAVIOR_CLEARED

Severity	informational
Description	<i>Component ID</i> no longer exhibits problematic behavior.
Troubleshooting	Please contact support.

DISK_PROBLEMATIC_BEHAVIOR_DETECTED

Severity	variable
Description	<i>Component ID</i> exhibits problematic behavior.
Troubleshooting	Please contact support.

DISK_PROBLEMATIC_BEHAVIOR_CLEARED

Severity	informational
Description	<i>Component ID</i> no longer exhibits problematic behavior.
Troubleshooting	Please contact support.

DISK_HIGH_MEDIA_ERROR_RATE_DETECTED

Severity	variable
Description	<i>Component ID</i> exhibits high media error rate of rule <i>rule_type</i> per cycle <i>cycle_type</i> .
Troubleshooting	Please contact support.

DISK_HIGH_MEDIA_ERROR_RATE_CLEARED

Severity	informational
Description	<i>Component ID</i> no longer exhibits high media error rate.
Troubleshooting	Please contact support.

SSD_PROBLEMATIC_BEHAVIOR_DETECTED

Severity	variable
Description	<i>Component ID</i> exhibits problematic behavior.
Troubleshooting	Please contact support.

SSD_PROBLEMATIC_BEHAVIOR_CLEARED

Severity	informational
Description	<i>Component ID</i> no longer exhibits problematic behavior.
Troubleshooting	Please contact support.

SSD_HIGH_MEDIA_ERROR_RATE_DETECTED

Severity	variable
Description	<i>Component ID</i> exhibits high media error rate of rule <i>rule_type</i> .
Troubleshooting	Please contact support.

SSD_HIGH_MEDIA_ERROR_RATE_CLEARED

Severity	informational
Description	<i>Component ID</i> no longer exhibits high media error rate.
Troubleshooting	Please contact support.

IB_SWITCH_FAILED

Severity	critical
Description	IB Switch <i>Component ID</i> failed.
Troubleshooting	Please contact support.

IB_SWITCH_STARTED_PHASEOUT

Severity	informational
Description	System started phasing out <i>Component ID</i> .
Troubleshooting	N/A

IB_SWITCH_STARTED_PHASEIN

Severity	informational
Description	System started phasing in <i>Component ID</i> .
Troubleshooting	N/A

IB_SWITCH_CONFIG_FAILED

Severity	warning
Description	<i>Component ID</i> could not be configured; <i>reason</i>
Troubleshooting	Please contact support

IB_SWITCH_IS_NOT_ANSWERING

Severity	warning
Description	<i>Component ID</i> is not answering
Troubleshooting	Please contact support

IB_SWITCH_BEGAN_ANSWERING

Severity	informational
Description	<i>Component ID</i> began answering after a period it did not

IB_SWITCH_HAS_POWER_FAILURE

Severity	warning
Description	Power supply # <i>PS Number</i> failure for <i>Component ID</i> .
Troubleshooting	Check power cables

IB_SWITCH_POWER_FAILURE

Severity	warning
Description	Power supply # <i>PS Number</i> failure for <i>Component ID</i> .
Troubleshooting	Check power cables

IB_SWITCH_POWER_RESTORED

Severity	informational
Description	Power supply # <i>PS Number</i> restored for <i>Component ID</i> .

IB_SWITCH_HAS_FAN_FAILURE

Severity	warning
Description	Fan # <i>fan</i> failure for <i>Component ID</i> .
Troubleshooting	Check fans, replace switch

IB_SWITCH_FAN_FAILURE

Severity	warning
Description	Fan # <i>fan</i> failure for <i>Component ID</i> .
Troubleshooting	Check fans, replace switch

IB_SWITCH_FAN_RESTORED

Severity	informational
Description	Fan # <i>fan</i> restored for <i>Component ID</i> .

IB_SWITCH_HAS_TEMPERATURE

Severity	variable
Description	<i>Component ID</i> - temperature of <i>temperature sesnsor</i> is <i>temperatureC</i> .
Troubleshooting	Please contact support.

IB_SWITCH_HIGH_TEMPERATURE

Severity	variable
Description	<i>Component ID</i> - temperature of <i>temperature sesnsor</i> is <i>temperatureC</i> .
Troubleshooting	Please contact support.

IB_SWITCH_HAS_VOLTAGE

Severity	variable
Description	Voltage # <i>voltage number</i> is <i>Voltage</i> Milli Volt and deviation from expected voltage is <i>Voltage Deviation</i> Milli Volt (which is <i>Voltage Deviation Percent%</i>) for <i>Component ID</i> .
Troubleshooting	Check voltage, replace switch

IB_SWITCH_VOLTAGE_PROBLEM

Severity	variable
Description	Voltage # <i>voltage number</i> is <i>Voltage</i> Milli Volt and deviation from expected voltage is <i>Voltage Deviation</i> Milli Volt (which is <i>Voltage Deviation Percent%</i>) for <i>Component ID</i> .
Troubleshooting	Check voltage, replace switch

IB_SWITCH_VOLTAGE_RESTORED

Severity	informational
Description	Voltage # <i>voltage number</i> was restored for <i>Component ID</i> .

IB_SWITCH_BAD_INDICATION

Severity	warning
Description	We got bad indication of type <i>Indication</i> for <i>Component ID</i> .
Troubleshooting	Check switch

IB_SWITCH_INTERN_ERROR

Severity	warning
Description	We got bad indication of type <i>Indication</i> for <i>Component ID</i> .
Troubleshooting	Check switch

IB_SWITCH_BAD_MANAGEMENT_WIRING

Severity	warning
Description	The wiring for IB switch management is incorrect. IB switch <i>Component ID</i> should be connected to module <i>Component ID</i>

IB_SWITCH_FIRMWARE_INCOMPATIBLE

Severity	warning
Description	The firmware version of ' <i>Component ID</i> ' is ' <i>New Version</i> '. It should be ' <i>Old Version</i> '.
Troubleshooting	None

IB_SWITCH_FIRMWARE_UPDATE_IN_PROG

Severity	warning
Description	Firmware version of ' <i>Component ID</i> ' is ' <i>New Version</i> '. It should be ' <i>Old Version</i> '. Firmware will be updated. It may take a while.
Troubleshooting	Wait for IB switch to complete initialization.

IB_SWITCH_FIRMWARE_UPDATED

Severity	informational
Description	The firmware version of ' <i>Component ID</i> ' was updated to ' <i>New Version</i> '.
Troubleshooting	None.

IB_SWITCH_PORT_POWER_UP_FAILED

Severity	minor
Description	Power up failed for port ' <i>Component ID</i> '
Troubleshooting	Please contact support

IB_SWITCH_LOG_COLLECT_SUCCESSFUL

Severity	informational
Description	Log collection for IB switch ' <i>switch_id</i> ' completed successfully. Log can be found in module ' <i>log_module</i> ' in the following directory: ' <i>log_location</i> '.

IB_SWITCH_LOG_COLLECT_FAILED

Severity	warning
Description	Log collection for IB switch ' <i>switch_id</i> ' failed. Failure reason: ' <i>failure_reason</i> '.

IB_SWITCH_PORT_ACTIVATION_FAILED

Severity	minor
Description	Power up failed for port ' <i>Component ID</i> '
Troubleshooting	Please contact support

IB_SWITCH_ALL_PORTS_DOWN

Severity	major
Description	All the ports of <i>Component ID</i> are down.
Troubleshooting	Please contact support

IB_SWITCH_SOME_PORTS_UP

Severity	informational
Description	Some ports of <i>Component ID</i> are up now.
Troubleshooting	N/A

IB_PORT_MOVED

Severity	informational
Description	Infinibind module port ' <i>module_port</i> ' moved from ' <i>from_port</i> ' to ' <i>to_port</i> '.

IB_SWITCH_CHANGE_DETECTED

Severity	major
Description	<i>Component ID</i> has been changed from GUID <i>old_guid</i> to GUID <i>new_guid</i> .

Troubleshooting	If IB_Switch replacement was intended, there is no problem.

MODULE_HAS_ACQUIRED_DHCP_ADDRESS

Severity	informational
Description	Module <i>Module ID</i> acquired DHCP address as part of the module equip process

UPS_IS_NOT_OK

Severity	major
Description	<i>Component ID</i> is currently not functioning, Reason: <i>Problem Code</i> .
Troubleshooting	Please contact support.

UPS_IS_OK

Severity	informational
Description	<i>Component ID</i> is currently functioning.
Troubleshooting	N/A

UPS_RACK_STATUS_CHANGE

Severity	informational
Description	Rack <i>Rack ID</i> UPSs are in the following states: UPS:1= <i>UPS 1 State</i> UPS:2= <i>UPS 2 State</i> UPS:3= <i>UPS 3 State</i> .

UPS_WAS_RECONFIGURED

Severity	informational
Description	New configuration was uploaded to <i>Component ID</i> .

UPS_WAS_NOT_RECONFIGURED

Severity	warning
Description	Unable to load new configuration to <i>Component ID</i> .

UPS_NEEDS_A_FIRMWARE_UPGRADE

Severity	informational
Description	UPS <i>Component ID</i> needs firmware upgrade

AOS_FILE_UPLOADED_TO_UPS

Severity	informational
Description	AOS bios file was uploaded to <i>Component ID</i> .

SUMX_FILE_UPLOADED_TO_UPS

Severity	informational
Description	SUMX application file was uploaded to <i>Component ID</i> .

UPS_SELF_TEST_HAS_STARTED

Severity	informational
Description	A UPS self-test was started on UPS <i>Component ID</i> .

UPS_SELF_TEST_WILL_BE_STARTED

Severity	informational
Description	About to start a UPS self-test on UPS <i>Component ID</i> .

UPS_SELF_TEST_HAS_PASSED

Severity	informational
Description	A UPS self-test has passed on UPS <i>Component ID</i> .

UPS_SELF_TEST_HAS_FAILED

Severity	major
Description	A UPS self-test has failed on UPS <i>Component ID</i> .

UPS_SELF_TEST_WAS_SKIPPED

Severity	warning
Description	A UPS self-test for <i>Component ID</i> has been skipped. <i>Reason</i>

UPS_SELF_TEST_IS_POSSIBLY_INACCURATE

Severity	informational
Description	A UPS self-test for <i>Component ID</i> might be inaccurate, capacity is <i>Battery Capacity</i> and is lower or equal to the minimum capacity <i>Minimum Capacity for Self Test</i> .

UPS_NEEDS_A_MANUAL_SELF_TEST

Severity	warning
Description	<i>Component ID</i> finished component test but it requires a manual self-test.
Troubleshooting	A component test should include self-testing, due to network link failure to the UPS a manual self-test is needed by waiting for full charge of the UPS and then pressing the UPS self-test button.

UPS_WAS_SUCCESSFULLY_UPGRADED

Severity	informational
Description	A UPS firmware upgrade has been successfully completed on UPS <i>Component ID</i> .

UPS_UPGRADE_HAS_FAILED

Severity	warning
Description	A UPS firmware upgrade has failed on UPS <i>Component ID</i> .

UPS_WILL_BE_CONFIGURED_NOW

Severity	informational
Description	UPS <i>Component ID</i> will be loaded with new configuration file due to <i>Configure Reason</i> .

UPS_CHANGE_DETECTED

Severity	informational
Description	<i>Component ID</i> has been changed from a serial of <i>old_serial</i> to <i>new_serial</i> .
Troubleshooting	If UPS replacement was intended, there is no problem.

UPS_BATTERY_CHANGE_DETECTED

Severity	informational
Description	Battery <i>Battery ID</i> of <i>Component ID</i> has been changed from a serial of <i>old_serial</i> to <i>new_serial</i> .
Troubleshooting	If battery replacement was intended, there is no problem.

UPS_SELF_TEST_TIMED_OUT

Severity	warning
Description	<i>Component ID</i> self-test failed to complete in reasonable time.
Troubleshooting	N/A

UPS_CLOCK_OUT_OF_SYNC

Severity	minor
Description	<i>Component ID</i> clock is more than a day out of sync
Troubleshooting	If it repeats, contact support

UPS_POWER_LOAD_PERCENT_TOO_HIGH

Severity	major
Description	UPS <i>Component ID</i> has power load percent of <i>Power Load Percent</i> .
Troubleshooting	If it repeats, contact support

UPS_POWER_LOAD_PERCENT_OK

Severity	minor
Description	The power load percent of UPS <i>Component ID</i> is now OK.
Troubleshooting	N/A

UPS_REMAINING_RUNTIME_TOO_LOW

Severity	critical
Description	UPS <i>Component ID</i> has a remaining runtime of <i>Remaining Runtime</i> minutes, this value might not be enough for an emergency shutdown in case of a power failure.
Troubleshooting	Contact support

UPS_REMAINING_RUNTIME_OK

Severity	minor
Description	UPS <i>Component ID</i> has a remaining runtime of <i>Remaining Runtime</i> minutes, this value should be enough for an emergency shutdown in case of a power failure.
Troubleshooting	N/A

UPS_PREDICTIVE_REMAINING_RUNTIME_TOO_LOW

Severity	minor
Description	UPS <i>Component ID</i> has a remaining runtime of <i>Remaining Runtime</i> minutes, but the time will drop to <i>Predictive Remaining Runtime</i> minutes in case one of the other UPSes stops working.
Troubleshooting	Contact support

UPS_PREDICTIVE_REMAINING_RUNTIME_OK

Severity	informational
Description	UPS <i>Component ID</i> has a remaining runtime of <i>Remaining Runtime</i> minutes, the time will drop to <i>Predictive Remaining Runtime</i> minutes in case one of the other UPSes stops working.
Troubleshooting	Contact support

UPS_HAS_FAILED

Severity	critical
Description	<i>Component ID</i> failed.
Troubleshooting	Please contact support.

UPS_IS_ON_BATTERY

Severity	major
Description	<i>Component ID</i> switched to battery power.
Troubleshooting	Check power input.

UPS_ON_UTILITY_POWER

Severity	informational
Description	<i>Component ID</i> switched back to utility power.

UPS_IS_IN_BYPASS

Severity	major
Description	<i>Component ID</i> entered bypass state.
Troubleshooting	Please contact support.

UPS_OUT_OF_BYPASS

Severity	informational
Description	<i>Component ID</i> went out of bypass state.

UPS_IS_TURNED_OFF

Severity	critical
Description	<i>Component ID</i> was turned off and does not supply output power.

UPS_TURNED_ON

Severity	informational
Description	<i>Component ID</i> was turned on.

UPS_BATTERY_IS_WEAK

Severity	major
Description	<i>Component ID</i> has weak battery .
Troubleshooting	Replace battery.

UPS_BATTERY_IS_NOT_WEAK

Severity	major
Description	<i>Component ID</i> has no weak battery .

UPS_BATTERY_IS_PREDICTIVE_WEAK

Severity	major
Description	<i>Component ID</i> has weak battery .
Troubleshooting	Replace battery.

UPS_BATTERY_IS_NOT_PREDICTIVE_WEAK

Severity	major
Description	<i>Component ID</i> has no weak battery .

UPS_CONTROL_COMMAND_FAILED

Severity	variable
Description	UPS control command ' <i>Command</i> ' on <i>Component ID</i> failed.

ATS_STATUS_UNSTABLE

Severity	informational
Description	Two consecutive readouts from <i>Component ID</i> were not equal.
Troubleshooting	If this persists contact support.

ATS_BAD_REPLY

Severity	warning
Description	Invalid reply from ATS <i>Component ID</i> , possible ATS failure.
Troubleshooting	N/A

ATS_SET_INPUT_LINE

Severity	informational
Description	ATS <i>Component ID</i> input line will be set to <i>JInput Line</i> .
Troubleshooting	N/A

ATS_SET_INPUT_LINE_FAILURE

Severity	warning
Description	Failed setting ATS <i>Component ID</i> input line to <i>JInput Line</i> .
Troubleshooting	Try running command again, make sure input line wasn't switched by using <i>ats_list</i> and/or looking at the ATS LEDs. If all fail, switch the lines manually by disconnecting one input line and reconnecting it.

ATS_SET_INPUT_TO_DOWN_LINE

Severity	warning
Description	Line JInput Line of ATS Component ID is down, will not switch input line.
Troubleshooting	Make sure the target input line is up and functioning and try again.

ATS_NO_REDUNDANCY

Severity	warning
Description	Line JInput Line of ATS Component ID is down, there is no input line redundancy.
Troubleshooting	Contact support

UPS_HAS_TEMPERATURE

Severity	variable
Description	Component ID - temperature is <i>temperatureC</i> . Serial <i>serial</i> .
Troubleshooting	Please contact support.

UPS_CALIBRATION_PENDING

Severity	informational
Description	UPS calibration is pending for <i>Component ID</i> .

UPS_CALIBRATION_STARTED

Severity	informational
Description	UPS <i>Origin</i> calibration started for <i>Component ID</i> .

UPS_CALIBRATION_PASSED

Severity	informational
Description	UPS calibration passed for <i>Component ID</i> .

UPS_CALIBRATION_FAILED

Severity	major
Description	UPS calibration failed for <i>Component ID</i> .

UPS_CALIBRATION_CANCELED

Severity	informational
Description	UPS calibration canceled for <i>Component ID</i> , due to <i>Origin</i> cancel.

UPS_CALIBRATIONS_DISABLED

Severity	informational
Description	UPS calibrations will be disabled for the pre-defined blackout period to allow maintenance.

UPS_CALIBRATIONS_ENABLED

Severity	informational
Description	UPS automatic calibrations will be enabled.

UPS_LAST_CALIBRATION_DATE_IN_FUTURE

Severity	warning
Description	The last calibration date for <i>Component ID</i> is <i>Calibration Date</i> , which is in the future, assume it's time to calibrate it.

UPS_DISABLED

Severity	informational
Description	<i>Component ID</i> has been disabled (will not be monitored).

UPS_ENABLED

Severity	informational

Description	<i>Component ID</i> has been enabled.

UPS_NETWORK_LINK_UP

Severity	warning
Description	Network link to <i>Component ID</i> was regained.
Troubleshooting	N/A

UPS_NETWORK_LINK_IS_DOWN

Severity	critical
Description	Network link to UPS <i>Component ID</i> is down.
Troubleshooting	Please contact support.

UPS_SERIAL_LINK_UP

Severity	warning
Description	Serial link to <i>Component ID</i> was regained.
Troubleshooting	N/A

UPS_SERIAL_LINK_DOWN

Severity	warning
Description	Serial link to UPS <i>Component ID</i> is down.
Troubleshooting	Please contact support.

UPS_LINK_READ_IGNORED_DATA_TIMEDOUT

Severity	warning
Description	Link to UPS <i>Component ID</i> timed out skipping ignored data.
Troubleshooting	Please contact support.

UPS_CABLE_CHECK_FAILED

Severity	minor
Description	Cable check of <i>Component ID</i> failed, either its serial or network cables are crosswired with <i>Hostname</i> .
Troubleshooting	Either the cable is disconnected or network and serial cables are miswired.

UPS_CABLE_CHECK_PASSED

Severity	informational
Description	Cable check of <i>Component ID</i> passed.

UPS_CABLE_CHECK_CONNECTION_FAILED

Severity	minor
Description	Cable check of <i>Component ID</i> failed, cannot connect to <i>Failed Link</i> link.
Troubleshooting	Failed to access the UPS through serial or network link, contact support.

ATS_LINK_UP

Severity	informational
Description	Link to ATS <i>Component ID</i> was regained.
Troubleshooting	N/A

ATS_LINK_DOWN

Severity	warning
Description	Link to ATS <i>Component ID</i> is down.
Troubleshooting	Please contact support.

ATS_LINE_INPUT_IS_OFF

Severity	major

Description	ATS <i>Component ID</i> input line <i>JLine</i> turned off.
Troubleshooting	N/A

ATS_LINE_INPUT_IS_ON

Severity	informational
Description	ATS <i>Component ID</i> input line <i>JLine</i> turned on.
Troubleshooting	N/A

ATS_LINE_OUTPUT_OFF

Severity	major
Description	ATS <i>Component ID</i> output line <i>Name</i> turned off.
Troubleshooting	Please contact support.

ATS_LINE_OUTPUT_ON

Severity	informational
Description	ATS <i>Component ID</i> output line <i>Name</i> turned on.
Troubleshooting	N/A

ATS_SOURCE_HAS_SWITCHED

Severity	informational
Description	ATS <i>Component ID</i> source line switched from <i>JPrevious</i> to <i>JCurrent</i> .
Troubleshooting	N/A

ATS_SOURCE_HAS_SWITCHED_FOR_OUTLET

Severity	informational
Description	ATS <i>Component ID</i> source line for outlet <i>POutlet</i> switched from <i>JPrevious</i> to <i>JCurrent</i> .
Troubleshooting	N/A

ATS_INPUT_LINE_OVER_CURRENT

Severity	warning
Description	Over-Current on ATS <i>Component ID</i> input line <i>JLine</i> , <i>Phase</i> .
Troubleshooting	Please contact support.

ATS_INPUT_LINE_CURRENT_RECOVERY

Severity	informational
Description	Recovered from over-current on ATS <i>Component ID</i> input line <i>JLine</i> , <i>Phase</i> .
Troubleshooting	N/A

ATS_OUTLET_OVER_CURRENT

Severity	warning
Description	Over-Current on ATS <i>Component ID</i> outlet <i>POutlet</i> .
Troubleshooting	Please contact support.

ATS_OUTLET_CURRENT_RECOVERY

Severity	informational
Description	Recovered from over-current on ATS <i>Component ID</i> outlet <i>POutlet</i> .
Troubleshooting	N/A

ATS_FAILED

Severity	minor
Description	ATS <i>Component ID</i> exhibits a failure state.
Troubleshooting	Please contact support.

ATS_RECOVERED

Severity	informational
Description	ATS <i>Component ID</i> exited from the failure state.
Troubleshooting	N/A

MODULE_NO_IP_CONNECTIVITY

Severity	warning
Description	There is no IP connectivity to failed <i>Component Id</i> .
Troubleshooting	Information sent to event center in case of module failure.

MODULE_NO_BMC_CONNECTIVITY

Severity	warning
Description	There is no BMC connectivity to failed <i>Component Id</i> .
Troubleshooting	Information sent to event center in case of module failure.

IB_HCA_DIAGNOSTICS

Severity	warning
Description	Diagnostics from IB HCA for <i>Component Id</i> .
Troubleshooting	Information sent to event center in case of module failure.

IB_HCA_DIAGNOSTICS_FAILED

Severity	warning
Description	Diagnostics for IB HCA card of <i>Component Id</i> failed due to <i>Failure Reason</i> .
Troubleshooting	Information sent to the event center in case of module failure.

MODULE_FAILED_WAS_POWERED_OFF

Severity	major
Description	The failed module <i>Failed module</i> has been powered off.
Troubleshooting	Information sent to event center in case of module failure.

MODULE_FAILED_WAS_NOT_POWERED_OFF

Severity	major
Description	The failed module <i>Failed module</i> has not been powered off as a failsafe due to <i>Failed IPMI module</i> not having IPMI set.
Troubleshooting	Information sent to event center in case of module failure.

MODULE_FAILED_COULD_NOT_BE_POWERED_OFF

Severity	major
Description	The failed module <i>Failed module</i> could not be powered off.
Troubleshooting	Information sent to event center in case of module failure.

MODULE_FAILED_SHOULD_BE_POWERED_OFF

Severity	major
Description	The failed module <i>Failed module</i> should be powered off based upon <i>Log String</i> .
Troubleshooting	Information sent to event center in case of module failure.

MODULE_SDR_INFO

Severity	major
Description	<i>Component ID</i> IPMI SDR info.

MODULE_SDR_EVENT_FAILURE

Severity	major
Description	<i>Component ID</i> IPMI SDR Event Error.

MODULE_SEL_LOG

Severity	warning
Description	<i>Component ID:[Index] Log string.</i> Raw event data is 'd0 d1 d2'.
Troubleshooting	Please contact support.

MODULE_FAILURE_DATA

Severity	critical
Description	<i>Component ID: Count lines Log string.</i>
Troubleshooting	Please contact support.

CONNECTION_TO_MAINTENANCE_MODULE_IS_OK

Severity	informational
Description	The Maintenance module can now be reached from <i>Component ID</i> .
Troubleshooting	Please contact support

NO_CONNECTION_TO_MAINTENANCE_MODULE

Severity	major
Description	The Maintenance module can not be reached from <i>Component ID</i> .
Troubleshooting	Please contact support

MODULE_TEMPERATURE_INCONSISTENT_WITH_OTHERS

Severity	warning
Description	<i>Component ID</i> external temperature not consistent with other modules.

SYSTEM_TEMPERATURE_IS_ABOVE_NORMAL

Severity	warning
Description	System temperature is <i>System TemperatureC</i> , which is above the normal temperature.
Troubleshooting	Cool the system down.

SYSTEM_TEMPERATURE_IS_HIGH

Severity	minor
Description	System temperature is <i>System TemperatureC</i> , which is high.
Troubleshooting	Cool the system down.

SYSTEM_TEMPERATURE_IS_HIGH_AND_STABILIZING

Severity	warning
Description	System temperature is <i>System TemperatureC</i> . It is stabilizing, but still close to the maximal allowable value.
Troubleshooting	Cool the system down.

SYSTEM_TEMPERATURE_IS_CRITICALLY_HIGH

Severity	critical
Description	System temperature is <i>System TemperatureC</i> , which exceeds operational level. Please initiate shutdown sequence. Without further action, the system will automatically shut itself down if it reaches <i>Shutdown ThresholdC</i> .
Troubleshooting	Cool the system down immediately or shut down the system using 'shutdown -y' and contact support.

SYSTEM_TEMPERATURE_IS_CRITICALLY_HIGH_SHUTTING_DOWN

Severity	critical
Description	System temperature is <i>System TemperatureC</i> , which is critically high. Shutting down the system.
Troubleshooting	Shut down the system using 'shutdown -y' and contact support.

SYSTEM_TEMPERATURE_IS_CRITICALLY_HIGH_SHUTDOWN_IMMEDIATELY

Severity	critical
Description	System temperature is <i>System TemperatureC</i> - which is critically high - but automatic shutdown is disabled. Shut down the system immediately!
Troubleshooting	Cool the system down immediately or shut down the system using 'shutdown -y' and contact support.

SYSTEM_TEMPERATURE_IS_CRITICALLY_HIGH_SHUT_IT_DOWN_IMMEDIATELY

Severity	critical
Description	System temperature is <i>System TemperatureC</i> - which is critically high - but automatic shutdown is disabled. You need to manually shut down the system immediately!
Troubleshooting	Cool the system down immediately or shut down the system using 'shutdown -y' and contact support.

SYSTEM_TEMPERATURE_IS_TOO_HIGH_AND_STABILIZING

Severity	major
Description	System temperature is <i>System TemperatureC</i> . It is stabilizing, but is still higher than the maximal allowable value. If the system doesn't cool down soon, it might be automatically shut down.
Troubleshooting	Contact support.

SYSTEM_TEMPERATURE_IS_TOO_HIGH

Severity	major
Description	System temperature is <i>System TemperatureC</i> . It approaches the maximal allowable value.
Troubleshooting	Cool the system down and contact support.

SYSTEM_TEMPERATURE_TOO_HIGH_AND_STABILIZING

Severity	major
Description	System temperature is <i>System TemperatureC</i> . It is stabilizing, but is still higher than the maximal allowable value.
Troubleshooting	Cool the system down and contact support.

SYSTEM_TEMPERATURE_TOO_HIGH

Severity	major
Description	System temperature is <i>System TemperatureC</i> , which is higher than the maximal allowable value.
Troubleshooting	Cool the system down and contact support.

SYSTEM_TEMPERATURE_TOO_HIGH_AND_STABILIZING_SHUTDOWN

Severity	major
Description	System temperature is <i>System TemperatureC</i> . It is stabilizing, but is still higher than the maximal allowable value. If the system doesn't cool down soon, it will be automatically shut down.
Troubleshooting	Cool the system down and contact support.

SYSTEM_TEMPERATURE_TOO_HIGH_SHUTDOWN

Severity	major
Description	System temperature is <i>System TemperatureC</i> , which is higher than the maximal allowable value. If the system doesn't cool down soon, it will be automatically shut down.
Troubleshooting	Cool the system down and contact support.

SYSTEM_TEMPERATURE_IS_TOO_LOW

Severity	major
Description	System temperature is <i>System TemperatureC</i> , which is lower than the minimal allowable value.
Troubleshooting	Contact support.

SYSTEM_TEMPERATURE_IS_OK_NOW

Severity	informational
Description	System temperature is <i>System TemperatureC</i> , which is within allowed limits.
Troubleshooting	N/A

SYSTEM_TEMPERATURE_RISES_SUSPICIOUSLY_FAST

Severity	warning
Description	System temperature (<i>System Temperature C</i>) is rising suspiciously fast (from <i>Previous TemperatureC</i>). Check air conditioning system.
Troubleshooting	Check air conditioning system or contact support.

SERVICE_MODE_OF_SYSTEM_HAS_CHANGED

Severity	variable
Description	Service mode of system has changed. The urgency of service actions has changed from <i>Previous Maintenance Urgency</i> to <i>Maintenance Urgency</i> .

MM_OK

Severity	informational
Description	<i>Component ID</i> is now OK.
Troubleshooting	Please contact support.

MM_FAILED

Severity	major
Description	<i>Component ID</i> has failed. Hardware status: <i>Status</i> .
Troubleshooting	Please contact support.

MM_CHANGE_DETECTED

Severity	major
Description	<i>Component ID</i> has been changed from a serial of <i>old_serial</i> to <i>new_serial</i> .
Troubleshooting	Was this MM actually replaced?

MODULE_MCH_DATA

Severity	critical
Description	<i>Component ID</i> : <i>[Bus.Device.Function]</i> Count lines <i>Log string</i> .
Troubleshooting	Please contact support.

UPS_BAD_BATTERY_FOUND_IN_SELF_TEST

Severity	major
Description	<i>Component ID</i> has low capacity or discharge state following self test.
Troubleshooting	Replace battery.

UPS_TURNED_OFF_AFTER_SELF_TEST_BAD_BATTERY

Severity	major
Description	<i>Component ID</i> turned itself off as a result of a self test, the battery is bad.
Troubleshooting	Replace battery.

UPS_TURNED_OFF_DURING_CALIBRATION_BAD_BATTERY

Severity	major
Description	<i>Component ID</i> turned itself off during a calibration, the battery is bad.
Troubleshooting	Replace battery.

UPS_BAD_BATTERY_FOUND_IN_CALIBRATION

Severity	major
Description	<i>Component ID</i> is off after calibration.
Troubleshooting	Replace battery.

UPS_INITIALIZATION_TIMED_OUT

Severity	informational
Description	<i>Component ID</i> Initialization timed out.

IPMI_SEL_ENTRY_INFO

Severity	informational
Description	<i>Entry Name</i> SEL entry on <i>component ID</i> <i>Date Time</i> with data <i>d0=d0 d1=d1 d2=d2 dir=direction</i> .
Troubleshooting	Please contact support.

MEMORY_COMMITMENT_IS_NEAR_LIMIT

Severity	warning
Description	<i>module</i> is <i>difference</i> KB below memory commit limit - a low margin.
Troubleshooting	Please contact support

MEMORY_COMMITMENT_OK

Severity	informational
Description	<i>module</i> is <i>difference</i> KB below memory commit limit - returned to a safe margin.
Troubleshooting	There is no problem at the moment.

HAS_TOO_MANY_PROCESSES

Severity	critical
Description	<i>module</i> has <i>processes</i> processes running.
Troubleshooting	Please contact support

DISK_DOES_NOT_EXIST

Severity	major
Description	<i>Component ID</i> doesn't exist.
Troubleshooting	Please contact support.

DISK_FIRMWARE_NOT_UPGRADEABLE

Severity	major
Description	Firmware upgrade for <i>Component ID</i> failed.
Troubleshooting	Please contact support.

ENCRYPT_ENABLE_DRIVE_FAILED

Severity	major
Description	Failed to enable encryption for <i>Component ID</i> . Error code: <i>Failure Reason</i> .
Troubleshooting	Please contact support.

ENCRYPTING_DISK_CRYPTO_ERASE_FAILED

Severity	major
Description	Failed to cryptographically erase <i>Component ID</i> . Error code: <i>Failure Reason</i> .
Troubleshooting	Please contact support.

ENCRYPT_SECURITY_HEALTH_CHECK_DISK_DEGRADED

Severity	major
Description	<i>Component ID</i> failed periodic security health check, reason: <i>Reason</i> .
Troubleshooting	Please contact support.

COMPONENT_FRU_REJECTED

Severity	major
Description	<i>Component ID</i> - Failed FRU validation.
Troubleshooting	Please contact support.

COMPONENT_FRU_ACCEPTED_IMPLICIT

Severity	informational
Description	<i>Component ID</i> - Passed implicit FRU validation.
Troubleshooting	There is no problem at the moment

DISK_FAILED_SHORT_STANDARD_TEST

Severity	major
Description	<i>Component ID</i> - Failed short standard test.

Troubleshooting	Please contact support.

DISK_CHANGE_WAS_DETECTED

Severity	informational
Description	<i>Component ID</i> has been changed from a <i>Old Vendor-Old Model</i> with a serial of <i>Old Serial</i> and with a firmware of <i>Old Firmware</i> to a <i>New Vendor-New Model</i> with a serial of <i>New Serial</i> and with a firmware of <i>New Firmware</i> .
Troubleshooting	Confirm that the disk replacement was intended.

DISK_FIRMWARE_CHANGE_WAS_DETECTED

Severity	informational
Description	The firmware of <i>Component ID</i> has changed from ' <i>Old Firmware</i> ' to ' <i>New Firmware</i> '.
Troubleshooting	Confirm that the firmware change was intended.

FAILURE_TO_RUN_DISK_SHORT_STANDARD_TEST

Severity	major
Description	<i>Component ID</i> - Failed short standard test.
Troubleshooting	Please contact support.

DISK_BMS_ERROR_DETECTED

Severity	warning
Description	<i>Component ID</i> - BMS error detected: <i>Sense Key/Additional Sense Code/Additional Sense Code Qualifier Sense Key - Sense Code (LBA: LBA)</i> .
Troubleshooting	Please contact support.

DISK_EXCESSIVE_BMS_ACTIVITY

Severity	variable
Description	<i>Component ID</i> exhibits excessive BMS activity, fill time is <i>Time to fill BMS log</i> minutes.
Troubleshooting	Please contact support.

SATA_SMART_STATUS_READING_FAILED

Severity	warning
Description	reading SMART attributes of Disk ID failed. SMART trip value=
Troubleshooting	Please contact support.

SATA_SMART_STATUS_READING_FAILURE

Severity	warning
Description	<i>Component ID</i> reading SMART attributes failed. SMART trip value=
Troubleshooting	Please contact support.

DISK_SMART_STATUS_BAD

Severity	major
Description	<i>Component ID</i> - SMART status: Bad.
Troubleshooting	Please contact support.

DISK_SMART_STATUS_GOOD

Severity	informational
Description	<i>Component ID</i> - SMART status: Good.
Troubleshooting	Please contact support.

DISK_SMART_READING_FAILED

Severity	warning
Description	<i>Component ID</i> - SMART reading failed.
Troubleshooting	Please contact support.

DISK_SMART_READING_OK

Severity	informational
Description	<i>Component ID</i> - SMART reading OK.
Troubleshooting	Please contact support.

DISK_BLOCK_SIZE_IS_INVALID

Severity	major
Description	<i>Component ID</i> was formatted with invalid block size of <i>Block Size</i> .
Troubleshooting	Please contact support.

DISK_ENCRYPTION_BAND_IS_INVALID

Severity	major
Description	The encryption configuration (banding) for <i>Component ID</i> is incorrect.
Troubleshooting	Please contact support.

DISK_ENCRYPTION_BAND_NOT_SET

Severity	major
Description	The encryption configuration (banding) for <i>Component ID</i> was not set.
Troubleshooting	Please contact support.

DISK_ENCRYPTION_ENROLLMENT_STATE_UNDETERMINED

Severity	minor
Description	The encryption enrollment state of <i>Component ID</i> cannot be determined.
Troubleshooting	Please contact support.

NON_ENCRYPTING_DISK_IN_ENCRYPTION_CAPABLE_SYSTEM

Severity	major
Description	<i>Component ID</i> does not support encryption.
Troubleshooting	Please contact support.

ENCRYPTING_DISK_DISCOVERY_FAILED

Severity	major
Description	The encryption state of <i>Component ID</i> could not be queried.
Troubleshooting	Please contact support.

ENCRYPTING_DISK_UNLOCK_FAILED

Severity	major
Description	Could not unlock encrypted disk <i>Component ID</i> .
Troubleshooting	Please contact support.

SES_STATUS_IS_ABNORMAL_NOW

Severity	warning
Description	On <i>module</i> SES component ' <i>Sensor Name</i> ' is in state ' <i>State</i> '.
Troubleshooting	Please contact support.

SES_STATUS_IS_NORMAL_NOW

Severity	informational
Description	On <i>module</i> SES component ' <i>Sensor Name</i> ' is in state ' <i>State</i> '.
Troubleshooting	N/A

SES_ALARM_IS_SIGNALING

Severity	warning
Description	On <i>module</i> alarm of type ' <i>Component type</i> ' is signaling.
Troubleshooting	Please contact support.

SES_ALARM_NO_LONGER_SIGNALING

Severity	informational
Description	On <i>module</i> alarm of type ' <i>Component type</i> ' is no longer signaling.
Troubleshooting	Please contact support.

SES_DOOR_LATCH_UNLOCKED

Severity	warning
Description	The door latch on module <i>module</i> is unlocked.
Troubleshooting	Please contact support.

SES_DOOR_LATCH_LOCKED

Severity	informational
Description	The door latch on module <i>module</i> is locked properly.
Troubleshooting	N/A

SES_PDB_FAILURE

Severity	major
Description	The PDB on module <i>module</i> has failed.
Troubleshooting	Please contact support.

SES_PDB_WARNING_ON

Severity	warning
Description	The PDB on module <i>module</i> displays some warning conditions.
Troubleshooting	Please contact support.

SES_PDB_IS_WORKING

Severity	informational
Description	The PDB module <i>module</i> is working now.
Troubleshooting	N/A

SES_BMC_ERROR_SIGNAL

Severity	warning
Description	The BMC on module <i>module</i> signals an error.
Troubleshooting	Please contact support.

SES_BMC_IS_OK_NOW

Severity	informational
Description	The BMC on module <i>module</i> is OK now.
Troubleshooting	N/A

SES_ALARM_LED_IS_SIGNALING

Severity	warning
Description	On <i>module</i> alarm of type ' <i>Component type (LED type)</i> ' is signaling.
Troubleshooting	Please contact support.

SES_ALARM_LED_NO_LONGER_SIGNALING

Severity	informational
Description	On <i>module</i> alarm of type ' <i>Component type (LED type)</i> ' is no longer signaling.
Troubleshooting	Please contact support.

SES_EXPANDER_RESET

Severity	warning
Description	The SES expander on <i>module</i> has been reset. Reset type was: <i>Reset type</i> .

SES_EXPANDER_RESET_COUNTER_CLEAR

Severity	informational
Description	The SES expander reset counter on <i>module</i> has reached it's maxiaml value and will be cleared.
Troubleshooting	N/A

IB_HCA_OVERHEATING

Severity	major
Description	Adapter <i>Module Component ID</i> is overheating and reached a temperature of <i>IB HCA Temperature</i> .
Troubleshooting	Please contact support.

CNA_OVERHEATING

Severity	major
Description	Adapter <i>CNA Component ID</i> is overheating and reached a temperature of <i>CNA Temperature</i> .
Troubleshooting	Please contact support.

COMPONENT_TEMPERATURE_IS_NORMAL

Severity	informational
Description	<i>Component ID</i> temperature is <i>temperatureC</i> . The temperature is normal..

COMPONENT_TEMPERATURE_IS_HIGH

Severity	warning
Description	<i>Component ID</i> temperature is <i>temperatureC</i> . The temperature is high.

COMPONENT_TEMPERATURE_IS_HIGH_AND_STABILIZING

Severity	warning
Description	<i>Component ID</i> temperature is <i>temperatureC</i> . The temperature is stabilizing, but still high.

COMPONENT_TEMPERATURE_IS_HIGH_AND_DROPPING

Severity	warning
Description	<i>Component ID</i> temperature is <i>temperatureC</i> . The temperature is dropping, but still high.

COMPONENT_TEMPERATURE_IS_ABNORMALLY_HIGH

Severity	minor
Description	<i>Component ID</i> temperature is <i>temperatureC</i> . The temperature is abnormally high.

COMPONENT_TEMPERATURE_IS_ABNORMALLY_HIGH_AND_STABILIZING

Severity	minor
Description	<i>Component ID</i> temperature is <i>temperatureC</i> . The temperature is stabilizing, but still abnormally high.

COMPONENT_TEMPERATURE_IS_ABNORMALLY_HIGH_AND_DROPPING

Severity	minor
Description	<i>Component ID</i> temperature is <i>temperatureC</i> . The temperature is dropping, but still abnormally high.

COMPONENT_TEMPERATURE_IS_VERY_HIGH

Severity	major
Description	<i>Component ID</i> temperature is <i>temperatureC</i> . The temperature is very high and may effect on component performance or even damage it.

COMPONENT_TEMPERATURE_IS_VERY_HIGH_AND_STABILIZING

Severity	major
Description	<i>Component ID</i> temperature is <i>temperatureC</i> . The temperature is stabilizing, but still very high.

COMPONENT_TEMPERATURE_IS_VERY_HIGH_AND_DROPPING

Severity	major
Description	<i>Component ID</i> temperature is <i>temperatureC</i> . The temperature is dropping, but still very high.

COMPONENT_TEMPERATURE_IS_EXTREMELY_HIGH

Severity	critical
Description	<i>Component ID</i> temperature is <i>temperatureC</i> . The temperature is extremely high. The component may immediately fail and permanent damage may occur.

COMPONENT_TEMPERATURE_IS_RISING

Severity	variable
Description	<i>Component ID</i> temperature is <i>temperatureC</i> . Temperature is rising.

COMPONENT_TEMPERATURE_IS_STABILIZING

Severity	variable
Description	<i>Component ID</i> temperature is <i>temperatureC</i> . Temperature is stabilizing.

COMPONENT_TEMPERATURE_IS_DROPPING

Severity	variable
Description	<i>Component ID</i> temperature is <i>temperatureC</i> . Temperature is dropping.

DISK_MEDIA_PRE_SCAN_ON

Severity	warning
Description	<i>Component ID</i> - Disk media pre scan is ON.
Troubleshooting	Please contact support.

DISK_MEDIA_PRE_SCAN_OFF

Severity	informational
Description	<i>Component ID</i> - Disk media pre scan is OFF.
Troubleshooting	Please contact support.

SES_FAN_HAS_SPEED_CHANGED

Severity	informational
Description	<i>Fan</i> speed changed from <i>Previous RPM</i> RPM to <i>New RPM</i> RPM.

SES_FAN_STATUS_OK

Severity	informational
Description	<i>Fan</i> is now OK.

SES_FAN_STATUS_BAD

Severity	warning
Description	<i>Fan</i> is failed or off.

SES_PSU_STATUS_HAS_CHANGED

Severity	variable
Description	<i>psu</i> changed state from <i>old_state</i> to <i>new state</i> .

SES_PSU_VOLTAGE_OUT_OF_RANGE

Severity	minor
Description	The <i>Voltage Type</i> DC voltage sensor of <i>PSU</i> shows <i>Voltage</i> which is not in the range of <i>Low Threshold-High Threshold</i> .
Troubleshooting	Please contact support.

SES_PSU_VOLTAGE_OK

Severity	informational
Description	<i>PSU Voltage Type</i> output DC voltage value is now <i>Voltage</i> , which is within the valid range.

SES_PSU_MONITORING_UNAVAILABLE

Severity	minor
Description	Can't monitor <i>PSU</i> , but it seems to supply power.
Troubleshooting	Please contact support.

SES_DOUBLE_PSU_FAILURE

Severity	major
Description	Both PSUs on <i>Module</i> report critical failures, this is probably because of a faulty PDB.
Troubleshooting	Please contact support.

SES_AMBIENT_TEMPERATURE_SEVERITY_CHANGED

Severity	variable
Description	Module <i>module sensor desc</i> reports (<i>readingC</i>) <i>state desc.</i>
Troubleshooting	Adjust ambient temperature at this module intake.

SES_TEMPERATURE_SEVERITY_CHANGED

Severity	variable
Description	Module <i>module sensor desc</i> reports (<i>readingC</i>) <i>state desc.</i>
Troubleshooting	Please contact support.

SES_TEMPERATURE_OK

Severity	informational
Description	Module <i>module sensor desc</i> reports (<i>readingC</i>) <i>state desc.</i>
Troubleshooting	N/A

INVALID_USM_VERSION_FOUND

Severity	minor
Description	The USM version on <i>Component</i> is invalid
Troubleshooting	Please contact support.

USM_VERSION_UNEXPECTED

Severity	warning
Description	USM on <i>Component</i> - unexpected code-level, found ' <i>Label</i> ' which is old and should be upgraded
Troubleshooting	Please contact support.

SES_VERSION_UNEXPECTED

Severity	warning
Description	SES on <i>Component</i> - unexpected enclosure management code-level, found <i>Major.Minor.Level.Level2</i> which is old and should be upgraded
Troubleshooting	Please contact support.

PDB_FIRMWARE_VERSION_UNEXPECTED

Severity	warning
Description	PDB firmware on <i>Component</i> - unexpected code-level, found <i>Major.Minor</i> which is old and should be upgraded
Troubleshooting	Please contact support.

PCM_FIRMWARE_VERSION_UNEXPECTED

Severity	warning
Description	PCM- <i>Id</i> firmware on <i>Component</i> - unexpected code-level, found <i>Major.Minor</i> which is old and should be upgraded
Troubleshooting	Please contact support.

PSU_FIRMWARE_VERSION_UNEXPECTED

Severity	major
Description	<i>Component</i> , of model ' <i>Model</i> ', has an unexpected code-level <i>Major.Minor</i> , which is old and should be upgraded.
Troubleshooting	Please contact support.

INVALID_PSU_MODEL

Severity	major
Description	Invalid model ' <i>PSU Model</i> ' for <i>PSU</i> .
Troubleshooting	Replace PSU.

SUSPICIOUS_PSU_INFORMATION

Severity	warning
Description	Suspicious information was found for <i>PSU</i> , this might happen after a PSU replacement. Some of the hardware sensors monitoring will be disabled until the module is power cycled.
Troubleshooting	Was the PSU replaced? If yes, power cycle the module when possible, if not notify support.

PSU_MODEL_IS_OK_NOW

Severity	informational
Description	Model ' <i>PSU Model</i> ' for <i>PSU</i> is valid.
Troubleshooting	N/A

ILLEGAL_PSU_MODEL_MIX

Severity	major
Description	<i>PSU-1</i> , of model ' <i>PSU-1 Model</i> ', can't be installed together with <i>PSU-2</i> of model ' <i>PSU-2 Model</i> '.
Troubleshooting	Replace one of the PSUs to appropriate model.

PSU_MODEL_MIX_IS_OK_NOW

Severity	informational
Description	<i>PSU-1</i> , of model ' <i>PSU-1 Model</i> ', is compatible with <i>PSU-2</i> of model ' <i>PSU-2 Model</i> '.
Troubleshooting	N/A

ILLEGAL_PSU_FIRMWARE_VERSIONS_MIX

Severity	major
Description	<i>PSU-1</i> , of model ' <i>PSU-1 Model</i> ', has firmware version <i>Major.Minor</i> which is not compatible with firmware version <i>Major.Minor</i> installed on <i>PSU-2</i> .
Troubleshooting	Install the latest PSU firmware version on both PSUs.

FAN_CONTROLLER_FIRMWARE_VERSION_UNEXPECTED

Severity	warning
Description	Fan controller firmware on <i>Component</i> - unexpected code-level, found <i>Major.Minor.Level.Build</i> (configuration <i>Major.Minor.Level</i>) which is old and should be upgraded
Troubleshooting	Please contact support.

IPMI_BMC_FIRMWARE_VERSION_UNEXPECTED

Severity	warning
Description	IPMI BMC firmware on <i>Module</i> - version <i>fmajor.fminor.faux</i> is old and should be upgraded
Troubleshooting	Please contact support.

BIOS_VERSION_IS_UNEXPECTED

Severity	warning
Description	BIOS on <i>Module</i> version <i>fmajor.fminor.faux</i> is old and should be upgraded
Troubleshooting	Please contact support.

FPGA_VERSION_IS_UNEXPECTED

Severity	warning
Description	FPGA on <i>Module</i> version <i>fmajor.fminor.faux</i> is old and should be upgraded
Troubleshooting	Please contact support.

INFINIBAND_HCA_VERSION_UNEXPECTED

Severity	warning
Description	Adapter <i>Module</i> version <i>Major.Minor.Build</i> is old and should be upgraded
Troubleshooting	Please contact support.

CNA_VERSION_UNEXPECTED

Severity	warning
Description	Adapter <i>CNA</i> version <i>Major.Minor.Build</i> is old and should be upgraded
Troubleshooting	Please contact support.

IPMI_BMC_IS_IN_UNEXPECTED_STATE

Severity	warning
Description	IPMI BMC firmware on <i>Module</i> is in unexpected state, possibly non-responsive
Troubleshooting	Please contact support.

IPMI_WATCHDOG_HAS_ERRORS

Severity	major
Description	IPMI watchdog on <i>Module</i> experienced command errors.
Troubleshooting	Please contact support.

SAS_VERSION_IS_UNEXPECTED

Severity	warning
Description	SAS Controller Firmware version on module <i>Module</i> version <i>actual</i> is old and should be upgraded
Troubleshooting	Please contact support.

SAS_VERSION_IS_INCONSISTENT

Severity	minor
Description	SAS Controller Firmware version on module <i>Module</i> version <i>actual</i> is inconsistent with persistent version <i>persistent</i> .
Troubleshooting	Please contact support.

SAS_CONTROLLER_FAULT

Severity	warning
Description	SAS Controller Firmware on <i>component ID</i> faulted with code <i>Fault Code</i>
Troubleshooting	Please contact support.

SAS_CONTROLLER_FAULT_CLEARED

Severity	informational
Description	SAS Controller Firmware on <i>component ID</i> recovered from its fault state.
Troubleshooting	Please contact support.

SAS_RESET_DETECTED

Severity	warning
Description	SAS Controller reset was detected on <i>component ID</i> total <i>Reset Count</i> times.
Troubleshooting	Please contact support.

IPMI_NOT_WORKING

Severity	major
Description	IPMI is not working on <i>Module</i> .
Troubleshooting	Please contact support.

IPMI_SEL_ENTRY_CRIT

Severity	critical
Description	<i>Entry Name</i> SEL entry on <i>component ID</i> <i>Date Time</i> with data d0= <i>d0</i> d1= <i>d1</i> d2= <i>d2</i> dir= <i>direction</i> .
Troubleshooting	Please contact support.

FC_PORT_HAS_UNEXPECTED_FIRMWARE

Severity	minor
Description	fc port <i>Component</i> firmware version is old and should be upgraded: <i>active firmware</i> instead of <i>correct firmware</i> .
Troubleshooting	Please contact support.

FC_LINK_ERROR_THRESH_EXCEEDED

Severity	major
Description	FC port <i>Component</i> counter threshold exceeded: <i>value</i> limit: <i>max</i>
Troubleshooting	Please contact support.

FC_MONITOR_ALERT

Severity	informational
Description	FC Port <i>Component</i> counter <i>counter</i> is out of range: now <i>value</i> range: <i>min max</i>
Troubleshooting	Please contact support.

KDB_HALT

Severity	major
Description	<i>Component ID</i> was stopped on KDB.
Troubleshooting	Please contact support.

KDB_LOG

Severity	major
Description	<i>Component ID</i> : <i>Line Count</i> lines <i>Log string</i>
Troubleshooting	Please contact support.

MCE_LOG

Severity	major
Description	<i>Component ID</i> : <i>Log string</i> .
Troubleshooting	Please contact support.

NETWORK_LINK_IS_NOW_DOWN

Severity	major
Description	Network interface <i>Interface Role</i> # <i>Interface Index</i> on <i>Component ID</i> - link disconnected.
Troubleshooting	Please contact support.

NETWORK_LINK_IS_NOW_UP

Severity	informational
Description	Network interface <i>Interface Role</i> # <i>Interface Index</i> on <i>Component ID</i> - link regained.
Troubleshooting	Please contact support.

FC_LINK_IS_NOW_DOWN

Severity	major
Description	FC port <i>Component</i> - link disconnected.
Troubleshooting	Please contact support.

FC_LINK_IS_NOW_UP

Severity	informational
Description	FC port <i>Component</i> - link regained.

NETWORK_LINK_NO_DATA

Severity	minor
Description	Network interface <i>Interface Role</i> # <i>Interface Index</i> on <i>Component ID</i> - link has no data flowing through for the last <i>Time Not flowing</i> seconds.
Troubleshooting	Please contact support.

NETWORK_LINK_NO_DATA_LONG

Severity	major
Description	Network interface <i>Interface Role</i> # <i>Interface Index</i> on <i>Component ID</i> - link has no data flowing through for the last <i>Time Not flowing</i> seconds.
Troubleshooting	Please contact support.

NETWORK_LINK_HAS_DATA

Severity	warning
Description	Network interface <i>Interface Role</i> # <i>Interface Index</i> on <i>Component ID</i> - link has data flowing through again.
Troubleshooting	Please contact support.

NETWORK_LINK_WAS_RESET_CONSECUTIVELY

Severity	major
Description	Network interface <i>Interface Role</i> # <i>Interface Index</i> on <i>Component ID</i> - link was reset consecutively .

NETWORK_LINK_PARTIAL_LOSS

Severity	variable
Description	Network interface <i>Interface Role</i> # <i>Interface Role Index</i> on <i>Component ID</i> has partial packet loss at a rate of <i>Packet Error Rate</i> .
Troubleshooting	Please contact support.

NETWORK_LINK_RETURNED_TO_NORMAL

Severity	informational
Description	Network interface <i>Interface Role</i> # <i>Interface Role Index</i> on <i>Component ID</i> no longer has partial packet loss.

COMPONENT_NETWORK_LINK_IS_DOWN

Severity	major
Description	Network interface to <i>Connected Component</i> on <i>Component ID</i> - link disconnected.
Troubleshooting	Please contact support.

COMPONENT_NETWORK_LINK_IS_UP

Severity	informational
Description	Network interface to component <i>Connected Component</i> on <i>Component ID</i> - link regained.
Troubleshooting	Please contact support.

MM_ETH_PORT_IS_DOWN

Severity	major
Description	Network interface <i>Component ID</i> to the maintenance module is down.
Troubleshooting	Please contact support.

MM_ETH_PORT_IS_UP

Severity	informational
Description	Network interface <i>Component ID</i> to the maintenance module is up.
Troubleshooting	Please contact support.

DAISY_CHAIN_IS_MISWIRED

Severity	major
Description	Daisy chain on module <i>Module</i> , device <i>Device name</i> is miswired.

Troubleshooting	Please contact support.

DAISY_CHAIN_IS_WIRED_OK

Severity	informational
Description	Daisy chain on module <i>Module</i> , device <i>Device name</i> is wired ok.
Troubleshooting	Please contact support.

DAISY_CHAIN_LINK_IS_DOWN

Severity	major
Description	Daisy chain link on module <i>Module</i> , device <i>Device name</i> is down.
Troubleshooting	Please contact support.

DAISY_CHAIN_LINK_IS_UP

Severity	informational
Description	Daisy chain link on module <i>Module</i> , device <i>Device name</i> is up.
Troubleshooting	Please contact support.

ETHERNET_INTERFACE_RESET

Severity	informational
Description	Network interface <i>Component ID</i> was reset.
Troubleshooting	N/A

NETWORK_LINK_FLOW_CONTROL_OFF

Severity	minor
Description	Network interface <i>Interface Role #Interface Role Index</i> on <i>Component ID</i> has flow control turned off.
Troubleshooting	N/A

NETWORK_LINK_FLOW_CONTROL_ON

Severity	informational
Description	Network interface <i>Interface Role #Interface Role Index</i> on <i>Component ID</i> has flow control turned on.
Troubleshooting	N/A

CF_PARTITION_INCONSISTENT

Severity	major
Description	Some of the partitions on the compact flash on <i>Component ID</i> have inconsistencies
Troubleshooting	The compact flash has inconsistencies on some of it's partitions.

UPS_WARMSTART

Severity	informational
Description	<i>Component ID</i> did a warmstart
Troubleshooting	Normally this is fine, if there is another issue at around the same time, it is worth mentioning this event to the support

CPU_LOCKUP

Severity	major
Description	<i>Component ID</i> had a lockup on CPU <i>#locked_cpu</i>
Troubleshooting	Please contact support.

FS_CORRUPTED

Severity	warning
Description	Filesystem <i>Device of Module</i> is corrupted
Troubleshooting	Please contact support.

FS_GOOD

Severity	informational
Description	Filesystem <i>Device of Module</i> is ok.
Troubleshooting	Please contact support.

SERIAL_CONSOLE_LINK_DOWN

Severity	warning
Description	Serial console link of <i>Target Module</i> read by <i>Source Module</i> is down, due to <i>Failure Reason</i> .
Troubleshooting	Please contact support.

SERIAL_CONSOLE_LINK_MISWIRED

Severity	warning
Description	Serial console link of <i>Target Module</i> read by <i>Source Module</i> is miswired, we expected module id <i>Target Module ID</i> but received <i>Received Module ID</i> .
Troubleshooting	Please contact support.

SERIAL_CONSOLE_LINK_CORRECT

Severity	informational
Description	Serial console link of <i>Target Module</i> read by <i>Source Module</i> is now correct.
Troubleshooting	Please contact support.

PROCESS_CORE_DUMP

Severity	warning
Description	Process <i>Process Name</i> with pid <i>Process ID</i> killed by signal <i>Signal</i> on module <i>Module Component ID</i> .

TOO_MANY_SYSLOG_MSGS

Severity	informational
Description	Syslog messages got updated more than once today on module <i>Module Component ID</i> .
Troubleshooting	There is no problem at the moment.

SAS_LINK_STATE_CHANGE

Severity	variable
Description	SAS link <i>SAS Type[ID]</i> on module <i>Module</i> changed state from <i>State</i> to <i>State</i> .
Troubleshooting	Please contact support.

SAS_LINK_SPEED_CHANGE

Severity	variable
Description	SAS link <i>SAS Type[ID]</i> on module <i>Module</i> speed changed from <i>Old Speed</i> to <i>New Speed</i> .
Troubleshooting	Please contact support.

SAS_LINK_ERRORS

Severity	variable
Description	SAS link <i>SAS Type[ID]</i> on module <i>Module</i> has too many errors, <i>Delta</i> since last sample.
Troubleshooting	Please contact support.

SAS_LINK_NO_MORE_ERRORS

Severity	informational
Description	SAS link <i>SAS Type[ID]</i> on module <i>Module</i> no longer has errors, <i>Delta</i> since last sample.
Troubleshooting	N/A

SAS_LINK_TOO_MANY_RESETS

Severity	major
Description	SAS link <i>SAS Type[ID]</i> on module <i>Module</i> had <i>Delta</i> resets, only <i>Allowed</i> resets are allowed. Disk <i>Disk</i> will be automatically phased out.
Troubleshooting	Please contact support.

SAS_LINK_TOO_MANY_RESETS_PHASEOUT_DISK

Severity	minor
Description	SAS link <i>SAS Type[ID]</i> on module <i>Module</i> had <i>Delta</i> resets, only <i>Allowed</i> resets are allowed. Please phase out disk <i>Disk</i> .
Troubleshooting	Phase out the disk and contact support.

ERROR_ON_SATA_HOST

Severity	major
Description	SATA device failure on <i>Module</i> .
Troubleshooting	Please contact support.

ERROR_ON_BOOT_DEVICE

Severity	major
Description	Boot device failure on module <i>Module</i> .
Troubleshooting	Please contact support.

ERROR_ON_AHCI_PORT

Severity	major
Description	AHCI device failure on module <i>Module</i> port <i>Port</i> .
Troubleshooting	Please contact support.

AHCI_PORT_RESCAN_REQUEST_WAS_SENT

Severity	major
Description	Rescan request for ahci port <i>Port</i> on module <i>Module</i> was sent.
Troubleshooting	Please contact support.

RESET_DETECTED_ON_AHCI_PORT

Severity	minor
Description	Hard reset detected on AHCI device, module <i>Module</i> , port <i>Port</i> .
Troubleshooting	Please contact support.

ERROR_ON_REMOTE_BOOT_DEVICE

Severity	major
Description	Boot device failure on module <i>Module</i> .
Troubleshooting	Please contact support.

CPU_LAPIC_IDLE

Severity	major
Description	<i>Component ID</i> is nearing NMI_CPU problem on CPU # <i>locked_cpu</i>
Troubleshooting	Please contact support.

FRU_PRODUCT_INFO_AREA_CHECKSUM_UNREADABLE

Severity	major
Description	FRU product information area checksum on module <i>Component ID</i> can't be read.
Troubleshooting	Please contact support.

DIMM_COMPLIANCE_CHECK_FAILED

Severity	major
Description	Installed DIMMs on module <i>Component ID</i> do not conform to the specification: <i>Failure reason</i>
Troubleshooting	Please contact support.

DIMM_COMPLIANCE_CHECK_DIMM_FAILED

Severity	major
Description	DIMM in slot <i>DIMM id</i> , part number ' <i>Part number</i> ', on module <i>Component ID</i> does not conform to the specification.

CPU_COMPLIANCE_CHECK_FAILED

Severity	major
Description	Installed processor on module <i>Component ID</i> does not conform to the specification.
Troubleshooting	Please contact support.

PCI_ADAPTER_COMPLIANCE_CHECK_FAILED

Severity	major
Description	Specification compliance check failed for PCI adapters on module <i>Component ID</i> .
Troubleshooting	Please contact support.

PCI_ADAPTER_COMPLIANCE_CHECK_ADAPTER_FAILED

Severity	major
Description	PCI adapter on module <i>Component ID</i> does not conform to the specification. Adapter details: ' <i>Description</i> '.
Troubleshooting	Please contact support.

INFINIBAND_HCA_COMPLIANCE_CHECK_FAILED

Severity	major
Description	Adapter <i>Component ID</i> does not conform to the specification.
Troubleshooting	Please contact support.

CNA_COMPLIANCE_CHECK_FAILED

Severity	major
Description	Adapter <i>Component ID</i> does not conform to the specification.
Troubleshooting	Please contact support.

PCIE_SLOT_COMPLIANCE_CHECK_FAILED

Severity	major
Description	PCIE slot on module <i>Component ID</i> does not conform to the specification.
Troubleshooting	Please contact support.

PCIE_SLOT_COMPLIANCE_CHECK_SLOT_FAILED

Severity	major
Description	PCIE slot <i>ID</i> on module <i>Component ID</i> does not conform to the specification.
Troubleshooting	Please contact support.

PCIE_SLOT_COMPLIANCE_CHECK_SPEED_FAILED

Severity	major
Description	PCIE slot <i>ID</i> on module <i>Component ID</i> does not conform to the speed specifications.
Troubleshooting	Please contact support.

CF_COMPLIANCE_CHECK_FAILED

Severity	major
Description	Installed compact flash card on module <i>Component ID</i> does not conform to the specification.
Troubleshooting	Please contact support.

CF_BIGGER_THAN_EXPECTED

Severity	warning
Description	Installed compact flash card on module <i>Component ID</i> has a size of <i>SizeMB</i> which is bigger than the expected size of <i>Spec SizeMB</i> .
Troubleshooting	Please contact support.

SSD_COMPLIANCE_CHECK_FAILED

Severity	major
Description	Installed SSD <i>Component ID</i> does not conform to the specification.
Troubleshooting	Please contact support.

SSD_DOES_NOT_EXIST

Severity	major
Description	SSD <i>Component ID</i> doesn't exist.
Troubleshooting	Please contact support.

SSD_FIRMWARE_UPGRADE_FAILED

Severity	major
Description	SSD <i>Component ID</i> failed firmware upgrade
Troubleshooting	Please contact support.

SSD_FIRMWARE_UPGRADE_SKIPPED

Severity	warning
Description	SSD <i>Component ID</i> has a new firmware version already installed, upgrade skipped.
Troubleshooting	N/A

SSD_BIGGER_THAN_EXPECTED

Severity	warning
Description	Installed SSD <i>Component ID</i> has a size of <i>SizeGB</i> which is bigger than the expected size of <i>Spec SizeGB</i> .
Troubleshooting	Please contact support.

SSD_SMALLER_THAN_EXPECTED

Severity	major
Description	Installed SSD <i>Component ID</i> has a size of <i>SizeGB</i> which is smaller than the expected size of <i>Spec SizeGB</i> .
Troubleshooting	Please contact support.

SSD_FOUND_UNEXPECTED

Severity	warning
Description	SSD <i>Component ID</i> was found while SSD Caching feature is disabled.
Troubleshooting	Please contact support.

SSD_SECURE_ERASE_FAILED

Severity	warning
Description	SSD <i>Component ID</i> secure erase failed.
Troubleshooting	Please contact support.

SSD_GENERIC_SUPPORT_USED

Severity	minor
Description	SSD <i>Component ID</i> using default smart attributes.
Troubleshooting	Please contact support.

SSD_SMART_READING_FAILED

Severity	major
Description	SSD <i>Component ID</i> failed reading smart attributes.
Troubleshooting	Please contact support.

SSD_NEAR_WEAROUT

Severity	variable
Description	SSD <i>Component ID</i> has bad SMART status. Attribute <i>Attribute (Attribute)</i> has value of <i>Value</i> .
Troubleshooting	Please contact support.

SSD_WORN_OUT

Severity	variable
Description	SSD <i>Component ID</i> has very bad SMART status and must be replaced. Attribute <i>Attribute (Attribute)</i> has value of <i>Value</i> .
Troubleshooting	Please contact support.

SSD_CYCLE_INFO

Severity	variable
Description	SSD <i>Component ID</i> passed <i>Cycles</i> cycles.
Troubleshooting	Please contact support.

SSD_LIFE_GAUGE

Severity	variable
Description	SSD <i>Component ID</i> smart attribute LIFE GAUGE exceeds a threshold. Value is <i>Value</i> .
Troubleshooting	Please contact support.

SSD_CHANGE_WAS_DETECTED

Severity	informational
Description	<i>Component ID</i> has been changed.
Troubleshooting	Confirm that the ssd replacement was intended.

SSD_FIRMWARE_CHANGE_WAS_DETECTED

Severity	informational
Description	The firmware of <i>Component ID</i> has changed from ' <i>Old Firmware</i> ' to ' <i>New Firmware</i> '.
Troubleshooting	Confirm that the SSD firmware change was intended.

SSD_SMART_ATTRIBUTE_THRESHOLD

Severity	variable
Description	SSD <i>Component ID</i> smart attribute <i>Smart attribute (Attribute)</i> exceeds a threshold. Value is <i>Value</i> .
Troubleshooting	Please contact support.

SSD_SPEED_HAS_CHANGED

Severity	major
Description	SSD <i>Component ID</i> speed has changed to <i>Speed</i> Gbps
Troubleshooting	Please contact support.

MODULE_DOWNLOAD_FAILED

Severity	minor
Description	Failure occured trying to download current version of the system to module <i>Module ID</i> , failure reason: <i>Reason</i> .
Troubleshooting	Please contact support.

MEMORY_ECC_ERRORS_DETECTED

Severity	warning
Description	Memory ECC errors were detected on <i>Module</i> .
Troubleshooting	Please contact support.

DIMM_CORRECTABLE_ERROR_DETECTED

Severity	variable
Description	Memory correctable ECC errors were detected on <i>Module</i> , <i>Count</i> errors on DIMM channel <i>Channel</i> , position <i>Position</i> .
Troubleshooting	Please contact support.

DIMM_UNCORRECTABLE_ERROR_DETECTED

Severity	variable
Description	Memory uncorrectable ECC errors were detected on <i>Module</i> , <i>Count</i> errors on DIMM channel <i>Channel</i> , position <i>Position</i> .

DIMM_ERRORS_PHASING_OUT_MODULE

Severity	major
Description	<i>Module</i> will be phased out as we detected too many DIMM errors there.
Troubleshooting	Please contact support.

PCIE_LINK_FATAL_ERROR

Severity	major
Description	A fatal error was detected on a PCIe link in <i>Module</i> .
Troubleshooting	Please contact support.

PCIE_LINK_ERROR

Severity	warning
Description	An error was detected on a PCIe link in <i>Module</i> .
Troubleshooting	Please contact support.

GEM_STATE_DUMPED

Severity	informational
Description	The GEM state on <i>Module</i> was dumped, reason: <i>Reason</i> .
Troubleshooting	N/A

DISK_WAS_TURNED_OFF

Severity	minor
Description	<i>Disk Component ID</i> was turned off.
Troubleshooting	Please contact your Administrator.

DISK_WAS_TURNED_ON

Severity	informational
Description	<i>Disk Component ID</i> was turned on.
Troubleshooting	N/A

DISK_GLIST_SIZE_TOO_HIGH

Severity	variable
Description	<i>Disk Component ID</i> GLIST size is <i>Glist Size</i> , which is too high.
Troubleshooting	Please contact support.

DISK_GLIST_CHANGED

Severity	warning
Description	<i>Disk Component ID</i> GLIST changed from <i>Previous glist</i> size to <i>Current glist</i> Size.
Troubleshooting	Please contact support.

DISK_IS_NOW_OFFLINE

Severity	warning
Description	<i>Disk Component ID</i> is now offline. It has been taken offline by the SCSI mid-layer.
Troubleshooting	Please contact support.

DISK_LOG_PAGE_READING_FAILED

Severity	warning
Description	<i>Disk Component ID</i> Failed reading log page. Opcode is <i>opcode</i> , page code is <i>page code</i> .
Troubleshooting	N/A

SSD_LOG_PAGE_READING_FAILED

Severity	warning
Description	SSD <i>Component ID</i> Failed reading log page <i>log</i> .
Troubleshooting	N/A

MODULE_IS_MISSING_MEMORY

Severity	major
Description	<i>Component ID</i> is missing memory. Actual memory size is <i>actual_mem</i> GB but should be <i>req_mem</i> GB.
Troubleshooting	Please contact your Administrator.

MODULE_IS_MISSING_REQUIRED_MEMORY

Severity	major
Description	<i>Component ID</i> has less memory (<i>actual_mem</i> GB) than is defined for use (<i>req_mem</i> GB).
Troubleshooting	Please contact your Administrator.

MODULE_HAS_MORE_MEMORY_THAN_EXPECTED

Severity	informational
Description	<i>Component ID</i> has more memory than expected. actual memory size is : <i>actual_mem</i> GB ,should be : <i>req_mem</i> GB.
Troubleshooting	Please contact your Administrator.

DISK_HIGH_READ_CORRECTED_WITH_DELAY_RATE

Severity	variable
Description	Disk <i>Component ID</i> has <i>number of read corrected with delay</i> read corrected errors with delay rate <i>rate</i> .
Troubleshooting	Please contact your Administrator.

DISK_FIRMWARE_VERSION_UNEXPECTED

Severity	warning
Description	Disks in module <i>Component</i> have unexpected firmware level. Check full event details for the complete list of the mismatched disks.
Troubleshooting	Please update the firmware level.

SSD_FIRMWARE_VERSION_UNEXPECTED

Severity	warning
Description	SSDs in module <i>Component</i> have unexpected firmware level. Check full event details for the complete list of the mismatched ssds.
Troubleshooting	Please update the firmware level.

MODULE_ROOT_FILESYSTEM_MIGHT_BE_CORRUPTED

Severity	minor
Description	The root file system of <i>Component ID</i> might be corrupted.
Troubleshooting	Please contact support.

INVALID_PKEY_DETECTED

Severity	warning
Description	Detected invalid IB PKEY in <i>Component ID</i> on port <i>Port</i> .

IPMI_USER_SETUP_FAILED

Severity	major
Description	Error setting up IPMI USER for module <i>Module</i> .
Troubleshooting	Contact support.

IPMI_LAN_SETUP_FAILED

Severity	major
Description	Error setting up IPMI LAN channel <i>Channel</i> for module <i>Module</i> .
Troubleshooting	Contact support.

FRU_PRODUCT_INFO_AREA_CHECKSUM_IS_INCORRECT

Severity	minor
Description	FRU product information area checksum on module <i>Component ID</i> is incorrect and can't be automatically fixed.
Troubleshooting	Please contact support.

FRU_PRODUCT_INFO_AREA_CHECKSUM_WAS_INCORRECT

Severity	warning
Description	FRU product information area checksum on module <i>Component ID</i> was incorrect and had to be reset.
Troubleshooting	N/A

FRU_PRODUCT_INFO_AREA_UNREADABLE

Severity	minor
Description	FRU product information area on module <i>Component ID</i> can't be read.
Troubleshooting	Please contact support.

IPMI_SEL_LOG_ENTRY

Severity	variable
Description	IPMI SEL log entry ' <i>Entry Name</i> ' found on <i>component ID</i> at ' <i>Date Time</i> '. Raw event data is ' <i>d0 d1 d2</i> ', direction is <i>direction</i> .
Troubleshooting	Please contact support.

IPMI_SEL_ENTRY_NEW

Severity	major
Description	<i>Entry Name</i> SEL entry on <i>component ID</i> <i>Date Time</i> <i>S</i> Type= <i>sensor_type</i> <i>S</i> Num= <i>sensor_number</i> with data <i>d0</i> = <i>d0</i> <i>d1</i> = <i>d1</i> <i>d2</i> = <i>d2</i> <i>dir</i> = <i>direction</i> .
Troubleshooting	Please contact support.

MSR_CPU_REG_INFO

Severity	major
Description	<i>Component ID</i> MSR CPU registers info.

INTERFACE_NODE_CHOKING

Severity	warning
Description	Interface node <i>#Node</i> has long Cache latency. Entered choked state <i>choke_time</i> msec ago on node= <i>Node</i>
Troubleshooting	Please contact support.

INTERFACE_NODE_UNCHOKING

Severity	variable
Description	Interface node <i>#Node</i> is leaving choked state after <i>choke_time</i> msec. Longest choke on node= <i>Node</i>
Troubleshooting	Please contact support.

INTERFACE_CHOKE_REPORT

Severity	informational
Description	Interface node <i>#Node</i> choked for a total of <i>choke_time</i> msecs in the last <i>report_period</i> minutes. The longest choke was of <i>longest_choke</i> msecs on node= <i>Node-ID</i>
Troubleshooting	Please contact support.

INTERFACE_ABORTS_REPORT

Severity	warning
Description	Interface node <i>#Node</i> handled <i>num_aborts</i> aborts in the last <i>report_period</i> seconds.
Troubleshooting	Please contact support.

VERY_LONG_LATENCY_TO_CACHE

Severity	variable
Description	Interface <i>#Node</i> has long latencies to the caches for <i>interval</i> seconds. Longest latency is <i>longest_latency</i> msecs to node <i>cache</i>
Troubleshooting	Please contact support.

LONG_LATENCY_TO_CACHE

Severity	variable
Description	Interface #Node had long latencies to the caches for <i>interval</i> msecs. Longest latency is <i>longest_latency</i> msecs to node <i>cache</i> .
Troubleshooting	Please contact support.

INODE_LONG_LATENCY_TO_COMPRESSION_NODE

Severity	variable
Description	Interface #Node had long latencies to the compressoin nodes for <i>interval</i> msecs. Longest latency is <i>longest_latency</i> msecs to node <i>compression</i> .
Troubleshooting	Please contact support.

INODE_VERY_LONG_LATENCY_TO_COMPRESSION_NODE

Severity	variable
Description	Interface #Node has long latencies to the compression nodes for <i>interval</i> seconds. Longest latency is <i>longest_latency</i> msecs to node <i>compression</i> .
Troubleshooting	Please contact support.

TASK_MANAGEMENT_FUNCTION_RECEIVED

Severity	minor
Description	Interface node #Node got task management function <i>task_management_function_code</i> from Host <i>host_name</i> at <i>port_type</i> port <i>port_name</i> to volume <i>volume</i> .
Troubleshooting	N/A

PERF_CLASS_RESOURE_EXHAUSTION

Severity	warning
Description	Exhausted all allowed resources for performance classes on <i>Module Id</i> , BUSY until resources available.

PERF_CLASS_RESOURCE_EXHAUSTION

Severity	warning
Description	Exhausted all allowed resources for performance classes on <i>Module Id</i> , BUSY until resources available.

CONNECTED_HOSTS_LIMIT_REACHED

Severity	informational
Description	Number of connected Hosts was reached for port ' <i>port_id</i> ' in Module <i>Module Id</i> .

QoS_HAS_BEEN_TRIGGERED

Severity	informational
Description	Queues on port ' <i>port_id</i> ' in Module <i>Module Id</i> caused QoS to be activated.

INODE_CONNECTED_TO_ALL_CACHES

Severity	informational
Description	Interface ' <i>node_id</i> ' is connected to all cache nodes.

IO_FAILURES_TO_ALL_CACHES

Severity	warning
Description	Interface node #Node got bad statuses from all caches during the last <i>failure_duration</i> seconds (last <i>failure_count</i> IOs) and therefore killed itself.

PERF_CLASS_RATE_AT_LIMIT

Severity	informational
Description	Performance class ' <i>perf_class</i> ' on <i>Module Id</i> reached its limit of <i>Limit Limit Name</i> , IOs being throttled.

PARTIAL_WRITE

Severity	warning
Description	Interface node <i>#Node</i> reports partial write to volume ' <i>volume</i> ' volume_uid= <i>volume_uid</i> , on lba <i>lba</i> , partition= <i>Logical Partition Number</i> , primary failed= <i>primary_failed</i> , secondary failed= <i>secondary_failed</i> , remote failed= <i>remote_failed</i> .

PORT_PREP_FOR_UPGRADE_TIMED_OUT

Severity	warning
Description	Preparation of <i>port_type</i> port ' <i>local_port_name</i> ' for hot-upgrade timed out due to host ' <i>host_name</i> ' port ' <i>host_port_name</i> ' host_port_addr

INTERFACE_DISCONNECTED_FROM_TARGET

Severity	major
Description	Interface node on module <i>module</i> cannot access target ' <i>target</i> ' through any gateway module.

INTERFACE_RECONNECTED_TO_TARGET

Severity	major
Description	Interface node on module <i>module</i> can access target ' <i>target</i> '.

ZOMBIE_TASK

Severity	warning
Description	Interface node <i>#Node</i> got a zombie task with op <i>Oxcmd</i> on volume ' <i>volume</i> '. total zombie tasks: <i>ztasks</i> total zombie pages: <i>zpages</i> .
Troubleshooting	N/A

METADATA_SERVICE_ENABLE

Severity	informational
Description	Metadata service is now enabled
Troubleshooting	N/A

METADATA_SERVICE_DB_CREATE

Severity	informational
Description	Database <i>DB</i> was created
Troubleshooting	N/A

METADATA_SERVICE_DB_DELETE

Severity	informational
Description	Database <i>DB</i> was deleted
Troubleshooting	N/A

IPINTERFACE_CREATE

Severity	informational
Description	A new iscsi IP interface was defined with name ' <i>interface name</i> ' on module <i>module</i> with ports ' <i>port list</i> ' and IP address <i>IP address</i>

IPINTERFACE_DELETE

Severity	informational
Description	ISCSI IP interface with name ' <i>interface name</i> ' was deleted

IPINTERFACE_RENAME

Severity	informational
Description	ISCSI IP interface with name ' <i>old name</i> ' and was renamed ' <i>interface name</i> '

IPINTERFACE_ADD_PORT

Severity	informational
Description	Port <i>#port index</i> was added to ISCSI IP interface with name ' <i>interface name</i> '

IPINTERFACE_REMOVE_PORT

Severity	informational
Description	Port # <i>port index</i> was removed from iSCSI IP interface with name ' <i>interface name</i> '

IPINTERFACE_UPDATE

Severity	informational
Description	iSCSI IP interface with name ' <i>interface name</i> ' was updated. Its IP address is <i>IP address</i>

IPINTERFACE_UPDATE_MANAGEMENT

Severity	informational
Description	Management IP interfaces were updated. Management IPs are <i>IP addresses</i>

IPINTERFACE_UPDATE_MANAGEMENT_IPV6

Severity	informational
Description	Management IP interfaces were updated. Management IPv6 addresses are <i>IPv6 addresses</i>

IPINTERFACE_UPDATE_VPN

Severity	informational
Description	VPN IP interfaces were updated. VPN IPs are <i>IP addresses</i>

IPINTERFACE_UPDATE_VPN_IPV6

Severity	informational
Description	VPN IPv6 interfaces were updated. VPN IPv6 addresses are <i>IP addresses</i>

AUXILIARY_INTERNAL_PORTS_ENABLED

Severity	informational
Description	<i>Port Count</i> auxiliary internal Ethernet ports were enabled

AUXILIARY_INTERNAL_PORTS_DISABLED

Severity	informational
Description	<i>Port Count</i> auxiliary internal Ethernet ports were disabled

IPSEC_ENABLED

Severity	informational
Description	IPSec was enabled

IPSEC_DISABLED

Severity	informational
Description	IPSec was disabled

IPSEC_CONNECTION_ADDED

Severity	informational
Description	A new IPSec connection named ' <i>name</i> ' was added

IPSEC_CONNECTION_UPDATED

Severity	informational
Description	The IPSec connection named ' <i>name</i> ' was updated

IPSEC_CONNECTION_REMOVED

Severity	informational
Description	The IPSec connection named ' <i>name</i> ' was removed

PRIVATE_KEY_ADDED

Severity	informational
Description	A new private key named ' <i>name</i> ' with fingerprint ' <i>fingerprint</i> ' and size <i>key_size</i> bits was added.

CERTIFICATE_REMOVED

Severity	informational
Description	The certificate named ' <i>name</i> ' was removed.

PKCS12_CERTIFICATE_ADDED

Severity	informational
Description	A new PKCS#12 named ' <i>name</i> ' with fingerprint ' <i>fingerprint</i> ' was added.

PKI_RENAME

Severity	informational
Description	PKI with the name ' <i>old name</i> ' was renamed to ' <i>new name</i> '

PKI_UPDATED

Severity	informational
Description	PKI with the name ' <i>name</i> ' and fingerprint ' <i>fingerprint</i> ' was updated

EMAIL_HAS_FAILED

Severity	variable
Description	Sending event <i>Event Code (Event Index)</i> to <i>Destination List</i> via <i>SMTP Gateway</i> failed. Module: <i>Module ID</i> ; Error message: ' <i>Error Message</i> '; timeout expired: <i>Timeout Expired?</i> .
Troubleshooting	Please contact support.

BULK_EMAIL_HAS_FAILED

Severity	variable
Description	Sending bulk email with <i>Events Number</i> events to <i>Destination List</i> via <i>SMTP Gateway</i> failed. Module: <i>Module ID</i> ; Error message: ' <i>Error Message</i> '; timeout expired: <i>Timeout Expired?</i> .
Troubleshooting	Please contact support.

SMS_HAS_FAILED

Severity	variable
Description	Sending event <i>Event Code (Event Index)</i> to <i>Destination List</i> via <i>SMS Gateway</i> and <i>SMTP Gateway</i> failed. Module: <i>Module ID</i> ; Error message: ' <i>Error Message</i> '; timeout expired: <i>Timeout Expired?</i> .
Troubleshooting	Please contact support.

HTTPS_HAS_FAILED

Severity	variable
Description	Sending event <i>Event Code (Event Index)</i> to <i>Destination List</i> via <i>HTTPS address</i> failed. Module: <i>Module ID</i> ; Error message: ' <i>Error Message</i> ' (<i>HTTP error code</i>); timeout expired: <i>Timeout Expired?</i> .
Troubleshooting	Please contact support.

EMAIL_NOT_SENT

Severity	variable
Description	Sending event <i>Event Code (Event Index)</i> to <i>Destination List</i> via <i>SMTP Gateway</i> was waived because of failed SMTP gateway. It will be not be used until <i>Retry Time</i> .
Troubleshooting	Please contact support.

SMS_NOT_SENT

Severity	variable
Description	Sending event <i>Event Code (Event Index)</i> to <i>Destination List</i> via <i>SMS Gateway</i> and <i>SMTP Gateway</i> was waived because of failed SMTP gateway. It will be not be used until <i>Retry Time</i> .

HEARTBEAT_EMAIL_HAS_FAILED

Severity	minor
Description	Sending heartbeat to <i>Destination Name</i> via <i>SMTP Gateway</i> failed. Module: <i>Module ID</i> ; Error message: ' <i>Error Message</i> '; timeout expired: <i>Timeout Expired?</i> .
Troubleshooting	Please contact support.

HEARTBEAT_SMS_HAS_FAILED

Severity	minor
Description	Sending heartbeat to <i>Destination Name</i> via <i>SMS Gateway</i> and <i>SMTP Gateway</i> failed. Module: <i>Module ID</i> ; Error message: ' <i>Error Message</i> '; timeout expired: <i>Timeout Expired?</i> .
Troubleshooting	Please contact support.

TEST_EMAIL_HAS_FAILED

Severity	minor
Description	Sending test to <i>Destination Name</i> via <i>SMTP Gateway</i> failed. Module: <i>Module ID</i> ; Error message: ' <i>Error Message</i> '; timeout expired: <i>Timeout Expired?</i> .
Troubleshooting	Please contact support.

TEST_SMS_HAS_FAILED

Severity	minor
Description	Sending test to <i>Destination Name</i> via <i>SMS Gateway</i> and <i>SMTP Gateway</i> failed. Module: <i>Module ID</i> ; Error message: ' <i>Error Message</i> '; timeout expired: <i>Timeout Expired?</i> .
Troubleshooting	Please contact support.

EVENTS_WERE_NOT_SAVED

Severity	variable
Description	Node <i>#Node</i> could not save <i>Num Lost</i> events of maximal severity <i>Maximal Severity</i> .
Troubleshooting	Please contact support.

EVENT_TEST

Severity	variable
Description	Event <i>Index</i> of <i>Count</i> test events (description: <i>Description</i> ; severity: <i>Severity</i> ; node: <i>Node</i>).
Troubleshooting	N/A

CUSTOM_EVENT

Severity	variable
Description	<i>Description</i>
Troubleshooting	N/A

INTERNAL_CUSTOM_EVENT

Severity	variable
Description	<i>Description</i>
Troubleshooting	N/A

CSS_PRODUCT_EVENT

Severity	variable
Description	CSS product ' <i>product</i> ' version ' <i>version</i> ' from server ' <i>server</i> ' running platform ' <i>platform</i> ' performed action ' <i>action</i> ' with properties ' <i>properties</i> '
Troubleshooting	N/A

MM_EVENT

Severity	variable
Description	Maintenance Module <i>Category</i> event : <i>Description</i>
Troubleshooting	Events coming from the Maintenance Module

TRACES_DUMPING

Severity	informational
Description	<i>Description</i>
Troubleshooting	N/A

SYSTEM_LOGS_COLLECTION

Severity	variable
Description	<i>Description</i>
Troubleshooting	N/A

SCRIPT_EXIT_STATUS

Severity	variable
Description	Script ' <i>Script Name</i> ' exited with status <i>Exit Status Description</i> .
Troubleshooting	N/A

USB_TO_SERIAL_CONNECTED_TO_WRONG_MODULE

Severity	warning
Description	The serial port <i>Serial Port</i> on module <i>Source Module</i> should be connected to <i>Target Module</i> but it is not.
Troubleshooting	Connect the serial cable to the correct module.

ALL_USB_TO_SERIAL_CONNECTED_CORRECTLY

Severity	informational
Description	All USB to serial cables are correctly connected to their designated modules.
Troubleshooting	N/A

FAILED_SAVING_EVENTS

Severity	warning
Description	Module <i>Module</i> failed saving events locally.
Troubleshooting	Please contact support.

XMPNS_ADMIN_CONTROL

Severity	informational
Description	<i>Action</i> ;username= <i>User</i>
Troubleshooting	N/A

XMPNS_USER_CONTROL

Severity	informational
Description	<i>Action</i> ;username= <i>User</i>
Troubleshooting	N/A

INTERCONNECT_TIMED_OUT

Severity	major
Description	Node # <i>Node ID</i> of type <i>Node Type</i> on <i>Module Component ID</i> could not establish connection to Node # <i>Node ID</i> of type <i>Node Type</i> on <i>Module Component ID</i> for <i>Timeout</i> seconds.
Troubleshooting	Please contact support.

EVENTS_WERE_LOST

Severity	variable
Description	<i>Number of events</i> Events of maximal severity ' <i>maximal severity</i> ' were lost due to overload.
Troubleshooting	Please contact support.

EVENTS_WERE_SUPPRESSED

Severity	variable

Description	<i>Number of Events Event Code</i> Events of maximal severity ' <i>maximal severity</i> ' were suppressed in the last <i>period</i> minutes.
Troubleshooting	N/A

NODE_HALTED

Severity	critical
Description	Node <i>#Node</i> has halted.
Troubleshooting	Please contact support

UPGRADE_MODULES_SW_NOT_IDENTICAL

Severity	warning
Description	File lists of different platform nodes are not identical to each other. Differing module is <i>differing_module</i>
Troubleshooting	Please contact support

UPGRADE_SOFTWARE_DOWNLOAD_FINISHED

Severity	informational
Description	Finished downloading software needed for upgrade to version <i>version</i> . Upgrade consequence is <i>consequence</i>
Troubleshooting	N/A

UPGRADE_FILE_LIST_RETRIEVAL_FAILED

Severity	critical
Description	Could not receive new version's file list from repository. Error code is <i>error</i> .
Troubleshooting	Contact support

UPGRADE_STARTS

Severity	informational
Description	System starting an upgrade.
Troubleshooting	N/A

PRE_UPGRADE

Severity	informational
Description	System preparing an upgrade procedure type <i>type</i> .
Troubleshooting	N/A

UPGRADE_IS_OVER

Severity	informational
Description	System went up after an upgrade.
Troubleshooting	N/A

IOS_RESTORED_AFTER_HOT_UPGRADE

Severity	informational
Description	System is able to perform I/Os after a hot upgrade.
Troubleshooting	N/A

UPGRADE_NO_NEW_FILES_FOR_UPGRADE

Severity	warning
Description	Repository version does not contain any new files. current version <i>current_version</i> new version is <i>new_version</i>
Troubleshooting	Contact support

UPGRADE_DOWNLOAD_REPOSITORY_COPY

Severity	critical
Description	Mirroring needed files from repository failed. Mirroring module is <i>mirroring_module</i> error is <i>error</i>
Troubleshooting	Contact support

UPGRADE_LOCAL_VERSION_DOWNLOAD_FAILED

Severity	critical
Description	Failure to distribute new software internally. Error code is <i>error</i> .
Troubleshooting	Contact support

UPGRADE_WAS_CANCELLED

Severity	informational
Description	Upgrade was cancelled with reason <i>reason</i> .
Troubleshooting	Contact support

HOT_UPGRADE_ABORTED

Severity	critical
Description	Hot upgrade aborted with reason <i>reason</i> .
Troubleshooting	Contact support

HOT_UPGRADE_HAS_FAILED

Severity	critical
Description	Hot upgrade failed while <i>erroneous_state</i> .
Troubleshooting	Contact support

PRE_UPGRADE_SCRIPT_INVOCATION_FAILED

Severity	critical
Description	Invocation of pre-upgrade script failed with error <i>error</i> .
Troubleshooting	Contact support

POST_UPGRADE_SCRIPT_INVOCATION_FAILED

Severity	critical
Description	Invocation of post-upgrade script failed with error <i>error</i> .
Troubleshooting	Contact support

UPGRADE_IS_NOT_ALLOWED

Severity	critical
Description	One or more of the pre-upgrade validations failed.
Troubleshooting	Fix the problems pointed out in previous events and revalidate.

PRE_UPGRADE_VALIDATION_FAILED

Severity	critical
Description	One of the pre-upgrade validations failed with status <i>error</i> .
Troubleshooting	Contact support

UPGRADE_IS_ALLOWED

Severity	informational
Description	All of the pre-upgrade validations passed successfully.
Troubleshooting	N/A

POST_UPGRADE_SCRIPT_STARTED

Severity	informational
Description	Post-upgrade script started.
Troubleshooting	N/A

POST_UPGRADE_SCRIPT_FINISHED

Severity	informational
Description	Post-upgrade script finished successfully.
Troubleshooting	N/A

PRE_UPGRADE_SCRIPT_DISAPPROVES

Severity	critical
Description	Upgrade cannot commence because some of the validations in the pre-upgrade script failed. Explanation: <i>explanation</i> .
Troubleshooting	Correct the system state according to the explanation and try again

POST_UPGRADE_SCRIPT_REPORTED_FAILURE

Severity	critical
Description	Post upgrade script reported failure. Script output: <i>explanation</i> .
Troubleshooting	Correct the system state according to the explanation and try again

FILE_TOO_LARGE

Severity	critical
Description	File <i>Filename</i> with size <i>Size</i> is too large for the current system limits, upgrades and installs can fail.
Troubleshooting	If file is not part of the system, delete it and regenerate file list. If it is part of the system, increase the limit.

SHOULD_BE_EMERGENCY_SHUTDOWN

Severity	critical
Description	An emergency shutdown has been detected, but UPS control is disabled. Shutdown reason: <i>Shutdown Reason</i> .
Troubleshooting	Solve the UPS problem immediately or shut the system down using 'shutdown -y' and contact support.

UPS_PROBLEM_CAUSING_MAINTENANCE_MODE

Severity	informational
Description	UPSeS state causing system to enter maintenance mode.
Troubleshooting	N/A

ADMINISTRATOR_PROCESS_FAILURE

Severity	warning
Description	Administrator process <i>Process ID</i> on <i>Module</i> failed with error <i>Error Code</i> . The last command was <i>Command Type</i> (ID: <i>Command ID</i> , in progress: <i>Command in Progress</i>).
Troubleshooting	Please contact support

ADMINISTRATOR_PROCESS_UNAVAILABLE

Severity	warning
Description	Command dispatch to administrator failed with timeout. Command was <i>Command Type</i> (ID: <i>Command ID</i>).
Troubleshooting	Please contact support

SSL_SETTINGS_CHANGED

Severity	informational
Description	Admin server SSL settings were changed. SSLv2 enabled: <i>SSLv2 Enabled</i> , cipher list: <i>Cipher List</i> . Change succeeded on <i>Module</i> ? <i>Change Succeeded</i> .
Troubleshooting	N/A

METADATA_SET

Severity	warning
Description	<i>Object type</i> with name ' <i>Object name</i> ' has new metadata value.
Troubleshooting	N/A

METADATA_DELETE

Severity	warning
Description	Metadata object deleted for <i>Object type</i> with name ' <i>Object name</i> '.
Troubleshooting	N/A

LOCAL_STORAGE_IS_CLEAR

Severity	informational
Description	Local storage is clear.
Troubleshooting	N/A

LOCAL_STORAGE_CLEAR_FAILED

Severity	warning
Description	Local storage clear failed on <i>module</i> .
Troubleshooting	N/A

PATCH_SCRIPT_ADDED

Severity	informational
Description	Added patch <i>Patch Name</i> .
Troubleshooting	Was patch supposed to have been added.

PATCH_SCRIPT_UPDATED

Severity	informational
Description	Updated patch <i>Patch Name</i> .
Troubleshooting	N/A

PATCH_SCRIPT_DELETED

Severity	informational
Description	Deleted patch <i>Patch Name</i> .
Troubleshooting	N/A

MODULE_FAILED_TO_FETCH_PATCH_SCRIPT

Severity	warning
Description	Module <i>Module</i> failed to fetch patch script <i>Patch Name</i> .
Troubleshooting	N/A

PATCH_SCRIPT_FAILED_TO_EXECUTE

Severity	informational
Description	Patch script <i>Patch Name</i> execution failed on module <i>Module</i>
Troubleshooting	N/A

PATCH_SCRIPT_EXECUTION_STARTED

Severity	informational
Description	Patch script <i>Patch Name</i> execution on module <i>Module</i> started with pid <i>Process ID</i>
Troubleshooting	N/A

PATCH_SCRIPT_EXECUTION_ENDED

Severity	informational
Description	Patch script <i>Patch Name</i> execution on module <i>Module</i> with pid <i>Process ID</i> ended with return code <i>Return Code</i>
Troubleshooting	N/A

DOMAIN_CREATED

Severity	informational
Description	Domain <i>domain_name</i> has been created.
Troubleshooting	N/A

DOMAIN_UPDATED

Severity	informational
Description	Domain <i>domain_name</i> has been updated.
Troubleshooting	N/A

DOMAIN_RENAMED

Severity	informational
Description	Domain <i>old_name</i> has been renamed to <i>domain_name</i> .
Troubleshooting	N/A

DOMAIN_DELETED

Severity	informational
Description	Domain <i>domain_name</i> has been deleted.
Troubleshooting	N/A

POOL_ADDED_TO_DOMAIN

Severity	informational
Description	Pool <i>pool_name</i> has been added to domain <i>domain_name</i> .
Troubleshooting	N/A

POOL_REMOVED_FROM_DOMAIN

Severity	informational
Description	Pool <i>pool_name</i> has been removed from domain <i>domain_name</i> .
Troubleshooting	N/A

POOL_MOVED_BETWEEN_DOMAINS

Severity	informational
Description	Pool <i>pool_name</i> has been moved from domain <i>domain_name</i> to domain <i>domain_name</i> .
Troubleshooting	N/A

GROUPED_POOL_MOVED_BETWEEN_DOMAINS

Severity	informational
Description	Grouped Pool <i>gp_name</i> has been moved from domain <i>domain_name</i> to domain <i>domain_name</i> .
Troubleshooting	N/A

DOMAINS_AUTO_SHIFT_RESOURCES

Severity	informational
Description	Resources from domain <i>domain_name</i> to domain <i>domain_name</i> have been auto shifted.
Troubleshooting	N/A

OBJECT_ATTACHED_TO_DOMAIN

Severity	informational
Description	Object <i>object_name</i> of type <i>object_type</i> has been added to domain <i>domain_name</i> .
Troubleshooting	N/A

OBJECT_REMOVED_FROM_DOMAIN

Severity	informational
Description	Object <i>object_name</i> of type <i>object_type</i> has been removed from domain <i>domain_name</i> .
Troubleshooting	N/A

DOMAIN_MANAGED_ATTRIBUTE_SET

Severity	informational
Description	Domain <i>domain_name</i> managed attribute was set to <i>managed_attribute</i> .
Troubleshooting	N/A

REMOTE_SUPPORT_CONNECTED

Severity	informational
Description	System connected to remote support center <i>Destination</i> .

UNABLE_TO_CONNECT_TO_REMOTE_SUPPORT

Severity	minor
Description	System is unable to connect to any remote support center.

REMOTE_SUPPORT_CONNECTION_LOST

Severity	variable
Description	Connection to remote support center <i>Destination</i> failed while the connection was in state <i>Disconnected Session State</i> .

REMOTE_SUPPORT_TIMEOUT

Severity	variable
Description	Connection to remote support center <i>Destination</i> timed out while the connection was in state <i>Disconnected Session State</i> .

REMOTE_SUPPORT_IMMINENT_TIMEOUT

Severity	minor
Description	System is about to disconnect busy connection to remote support center <i>Destination</i> .

REMOTE_SUPPORT_DEFINED

Severity	informational
Description	Defined remote support center <i>Name</i> with IP address <i>Address</i> and port <i>Port</i> .

REMOTE_SUPPORT_DELETED

Severity	informational
Description	Deleted remote support center <i>Name</i> .

REMOTE_SUPPORT_DISCONNECTED

Severity	variable
Description	System disconnected from remote support center <i>Destination</i> while the connection was in state <i>Disconnected Session State</i> .

REMOTE_SUPPORT_CLIENT_MOVED

Severity	informational
Description	The remote support client moved from <i>Old Module</i> to <i>New Module</i> .

REMOTE_SUPPORT_CLIENT_NO_AVAILABLE_MODULES

Severity	minor
Description	No live modules with <i>Port Type</i> ports are available to run the remote support client.

REMOTE_SUPPORT_AUTOMATIC_SESSION_SETUP_ATTEMPTED

Severity	informational
Description	System attempt to open remote support session automatically.

TIMEZONE_SET

Severity	informational
Description	Timezone of the system was set to <i>Timezone</i> .
Troubleshooting	N/A

TIME_SET

Severity	informational
Description	On <i>Previous Time and Date</i> date and time of the system were set to <i>Time and Date</i> .
Troubleshooting	If date and/or time setting was intended, there is no problem.

MANAGER_OPERATION_BLOCKED

Severity	critical
Description	The Manager is blocked.
Troubleshooting	Please escalate to IBM XIV support.

MANAGER_OPERATION_RESUMED

Severity	informational
Description	The Manager is no longer blocked.
Troubleshooting	N/A

ELICENSE_ACCEPTED

Severity	informational
Description	Electronic license was accepted by ' <i>Approver Name</i> '.
Troubleshooting	N/A

ELICENSE_VIOLATION

Severity	warning
Description	Latest version of the electronic license was not approved.
Troubleshooting	Please approve the electronic license.

AUDIT_ENABLED

Severity	informational
Description	CLI command auditing activated.
Troubleshooting	N/A

AUDIT_DISABLED

Severity	warning
Description	CLI command auditing deactivated.
Troubleshooting	N/A

MASTER_SM_CHOSEN

Severity	informational
Description	Subnet manager on module ' <i>Component ID</i> ' becomes master.
Troubleshooting	N/A

MODULE_IB_PORTS_DOWN

Severity	warning
Description	SM: all infiniband ports on module ' <i>Component ID</i> ' are down..
Troubleshooting	N/A

MODULE_SM_PRIO_CHANGED

Severity	informational
Description	SM: SM priority changed on module ' <i>Component ID</i> ' from ' <i>Priority</i> ' to ' <i>Priority</i> '. Change reason: ' <i>Change Reason</i> '.
Troubleshooting	N/A

MASTER_SM_RESCAN

Severity	informational
Description	Master SM on module ' <i>Component ID</i> ' has reloaded configuration and rescans network.
Troubleshooting	N/A

IB_MODULE_MISWIRE

Severity	warning
Description	Infiniband miswire: Port ' <i>Port Number</i> ' of module <i>Component ID</i> should be connected to ' <i>Component ID</i> '. Disallowed GUID ' <i>GUID of IB Device (Switch Port)</i> ' is connected instead.
Troubleshooting	N/A

IB_SWITCH_MISWIRE

Severity	warning
Description	Infiniband miswire: 'Switch ID' with GUID 'Switch GUID' should be connected to 'Switch Port'. Disallowed GUID 'GUID of IB Device (Component ID)' is connected instead.
Troubleshooting	N/A

IB_LINK_UP

Severity	informational
Description	Link on 'Component ID (Component ID)' is up.
Troubleshooting	N/A

IB_LINK_DOWN

Severity	warning
Description	Link on 'Component ID (Component ID)' is down.
Troubleshooting	N/A

IB_MISSING_LINK

Severity	warning
Description	Link on 'Component ID (Component ID)' is missing.
Troubleshooting	N/A

IB_MODULE_PORT_BAD_GUID_FORMAT

Severity	warning
Description	Module port#Port Number, connected to switch port 'Component ID' has GUID:'Port GUID' which is not xiv-GUID or its serial is not compatible with the machine.
Troubleshooting	N/A

IB_BAD_CONFIGURATION

Severity	major
Description	Bad IB configuration values found. Given values are 'IB component ID'. Max values are 'IB component ID'. Min values are 'IB component ID'.
Troubleshooting	N/A

IB_BAD_PORT_PERFORMANCE

Severity	variable
Description	Bad Port Performance: too many errors of type 'Counter Name' on 'Component ID'. Action taken: 'Action Taken' Recommendation: 'Recommendation'.
Troubleshooting	N/A

IB_PORT_SHUTDOWN

Severity	warning
Description	Port 'Component ID' is going to be shutdown. Reason: number of Shutdown Reason exceeded threshold.
Troubleshooting	N/A

IB_BAD_LINK_PERFORMANCE

Severity	warning
Description	Bad Link Performance: link width is 'Link Width' link speed is 'Link Speed' on 'Component ID'
Troubleshooting	N/A

IB_SWITCH_MALFUNCTION

Severity	major
Description	Found a malfunction in switch 'Component ID'. Will shutdown the switch when possible (when redundant link will be available).
Troubleshooting	N/A

IB_SWITCH_IS_MISSING

Severity	warning
Description	An Infiniband switch is missing ' <i>Component ID</i> '.
Troubleshooting	N/A

IB_MISSING_SWITCH_FOUND

Severity	informational
Description	An Infiniband switch ' <i>Component ID</i> ' which was missing is now found.
Troubleshooting	N/A

IB_PORT_FORCE_DOWN_SUCCESS

Severity	critical
Description	IB port ' <i>HCA Port</i> ' on module ' <i>Module</i> ' was powered off.
Troubleshooting	Contact Support.

IB_PORT_FORCE_DOWN_FAILED

Severity	critical
Description	Failed to power off IB port ' <i>HCA Port</i> ' on module ' <i>Module</i> '.
Troubleshooting	Contact Support.

IB_PORT_FORCE_UP_SUCCESS

Severity	informational
Description	IB port ' <i>HCA Port</i> ' on module ' <i>Module</i> ' was powered on.
Troubleshooting	Contact Support.

IB_PORT_FORCE_UP_FAILED

Severity	warning
Description	Failed to power on IB port ' <i>HCA Port</i> ' on module ' <i>Module</i> '.
Troubleshooting	Contact Support.

INFINIBAND_ERROR

Severity	major
Description	Failed to load infiniband stack on module <i>Module ID</i> . Error is: <i>Error description</i> .
Troubleshooting	Possible physical problem with module. Contact support.

INFINIBAND_PORT_ERROR

Severity	major
Description	Failed to bring up Infiniband port <i>HCA Port Number</i> on module <i>Module ID</i> . Error is: <i>Error description</i> .
Troubleshooting	Possible physical problem with module. Contact support.

COMPONENT_REQUIRES_SERVICING

Severity	major
Description	Component <i>Component ID</i> requires service: <i>Component Required Service</i> , due to: <i>Component Service Reason</i> . The urgency of this service is <i>Maintenance Urgency</i>

COMPONENT_REQUIRES_IMMEDIATE_SERVICING

Severity	major
Description	Component <i>Component ID</i> which previously had it's service deferred now requires immediate service: <i>Component Required Service</i> , due to: <i>Component Service Reason</i>

COMPONENT_REQUIRED_SERVICE_CLEARED

Severity	informational
Description	Component <i>Component ID</i> does NOT require service anymore

PERF_CLASS_MAX_IO_RATE_UPDATED

Severity	informational
Description	Performance Class <i>name</i> max IO rate was changed to <i>IO rate</i> IOPS

PERF_CLASS_MAX_BW_RATE_UPDATED

Severity	informational
Description	Performance Class <i>name</i> max BW rate was changed to <i>BW rate</i> MBps

PERF_CLASS_CREATE

Severity	informational
Description	Performance Class with name ' <i>name</i> ' was created

PERF_CLASS_DELETE

Severity	informational
Description	Performance Class with name ' <i>name</i> ' was deleted

PERF_CLASS_ADD_HOST

Severity	informational
Description	Host with name ' <i>host_name</i> ' was added to Performance Class with name ' <i>name</i> '

PERF_CLASS_REMOVE_HOST

Severity	informational
Description	Host with name ' <i>host_name</i> ' was removed from Performance Class with name ' <i>name</i> '

PERF_CLASS_ADD_POOL

Severity	informational
Description	Pool with name ' <i>pool.name</i> ' was added to Performance Class with name ' <i>pool.perf_class</i> '

PERF_CLASS_REMOVE_POOL

Severity	informational
Description	Pool with name ' <i>pool.name</i> ' was removed from Performance Class with name ' <i>name</i> '

PERF_CLASS_ADD_DOMAIN

Severity	informational
Description	Domain <i>domain_name</i> was added to Performance Class <i>name</i>

PERF_CLASS_REMOVE_DOMAIN

Severity	informational
Description	Domain <i>domain_name</i> was removed from Performance Class <i>name</i>

VOLUME_MODIFIED_DURING_IO_PAUSE

Severity	warning
Description	Volume ' <i>vol_name</i> ' of CG ' <i>cg_name</i> ' was modified during Pause IO with token ' <i>token</i> '
Troubleshooting	Retry after completing CG changes.

CONS_GROUP_MODIFIED_DURING_IO_PAUSE

Severity	warning
Description	CG ' <i>cg_name</i> ' was modified during Pause IO with token ' <i>token</i> '
Troubleshooting	Retry after completing CG changes.

IO_PAUSED_FOR_CONS_GROUP

Severity	Informational
Description	Pause IO on CG with name ' <i>cg_name</i> ' was started with <i>timeoutms</i> timeout . Token is ' <i>token</i> '.

IO_RESUMED_FOR_CONS_GROUP_EXPLICITLY

Severity	Informational
Description	Pause IO on CG with name ' <i>cg_name</i> ' and token ' <i>token</i> ' was resumed by user request.

IO_RESUMED_FOR_CONS_GROUP_AUTOMATICALLY

Severity	Informational
Description	Pause IO on CG with name ' <i>cg_name</i> ' and token ' <i>token</i> ' was resumed after snapgroup creation.

IO_RESUMED_FOR_CONS_GROUP_UPON_SYSTEM_ERROR

Severity	warning
Description	Pause IO on CG with name ' <i>cg_name</i> ' and token ' <i>token</i> ' was resumed after system error.

IO_RESUMED_FOR_CONS_GROUP_UPON_TIMEOUT_EXPIRATION

Severity	warning
Description	Pause IO on CG with name ' <i>cg_name</i> ' and token ' <i>token</i> ' was canceled after timeout.
Troubleshooting	Use longer timeout value or require less time for performing action.

ISYNC_RPC_TIMEOUT

Severity	major
Description	ISync RPC ' <i>rpc_uid</i> ' on volume ' <i>vol_uid</i> ' from initiator uid ' <i>init_uid</i> ' is delayed for more than <i>timeout</i> seconds.

HOST_PROFILE_SET

Severity	informational
Description	Host profile value has been set for host <i>host_name</i> .
Troubleshooting	N/A

HOST_PROFILE_SET_REJECTED

Severity	warning
Description	Host profile setting was rejected because the host attempted to set the profile of host <i>host_name</i> more than once in the last <i>min_set_interval</i> minutes.
Troubleshooting	N/A

HOST_PROFILE_CLEARED

Severity	informational
Description	Host profile value has been cleared for host <i>host_name</i> .
Troubleshooting	N/A

SYMON_COMPONENT_FAILURE_RECOMMENDED

Severity	informational
Description	Symptomatic monitoring recommends failing <i>Component</i> .
Troubleshooting	N/A

SYMON_COMPONENT_FAILURE_WOULD_HAVE_BEEN_RECOMMENDED

Severity	informational
Description	Symptomatic monitoring would have recommended failing <i>Component</i> .
Troubleshooting	N/A

SYMON_INODES_ARE_DISTRESSED

Severity	informational
Description	Symptomatic monitoring detected that some interface nodes are distressed.
Troubleshooting	N/A

MEDIUM_ERROR_RECOVERED

Severity	informational
Description	Medium error on volume= <i>Volume</i> , logical-partition= <i>Logical Partition Number</i> , offsetted-logical-partition= <i>Offsetted Logical Partition Number</i> was recovered.
Troubleshooting	N/A

MEDIUM_ERROR_NOT_RECOVERED

Severity	critical
Description	Medium error on volume= <i>Volume</i> , logical-partition= <i>Logical Partition Number</i> , offsetted-logical-partition= <i>Offsetted Logical Partition Number</i> could not be recovered due to <i>Reason</i> .
Troubleshooting	N/A

ALU_CREATE

Severity	informational
Description	ALU was defined with name ' <i>ALU name</i> ' associated with host ' <i>ALU host name</i> ' lun ' <i>ALU lun</i> '.

ALU_DELETE

Severity	informational
Description	ALU with name ' <i>ALU name</i> ' associated with host ' <i>ALU host name</i> ' lun ' <i>ALU lun</i> ' was deleted.

ALU_UNBOUND_ALL

Severity	informational
Description	All SLUs of ALU with name ' <i>ALU name</i> ' were unbound.

Return codes

This section contains descriptions of CLI 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

For information about how to use the IBM Hyper-Scale Manager REST Application Programming Interface (API), see the IBM Hyper-Scale Manager documentation.

- [IBM Hyper-Scale Manager documentation](#)

XIV Open API

The following information describes how to use the IBM® XIV® Software System Open Application Programming Interface (API).

- [XIV Open API overview](#)
The IBM XIV Storage System Open Application Programming Interface (API) is a nonproprietary storage-management client application.
- [CIM agent for the XIV system](#)
The CIM agent is preinstalled on the administrative module. The embedded CIM agent is enabled automatically and preconfigured. You can manage the XIV system from the CIM agent that is bundled with the administrative module.
- [XIV Open API component definitions](#)
The XIV Open API elements include schemas, classes, properties, methods, indications, associations, references, and qualifiers.
- [Functional diagrams of the CIM agent](#)
The functional diagrams of the CIM agent object model show specific functions that the CIM agent provides. The diagrams also illustrate the architecture of the CIM agent for the XIV system.

- [Conformance tests](#)
All SNIA official Certification Test Programs (CTP) and Microsoft System Center Virtual Machine Manager (SCVMM) Storage Automation tests are complete.

XIV Open API overview

The IBM® XIV Storage System Open Application Programming Interface (API) is a nonproprietary storage-management client application.

The XIV 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 XIV 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 the XIV system. The XIV Open API presents another option for managing your systems by complementing the use of the IBM XIV Storage Management GUI and the IBM XIV command-line interface (XCLI). The XIV Open API is an embedded component of the XIV system.

You can implement the XIV 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 XIV Open API uses the CIM technology to manage proprietary devices as open system devices through storage management applications. The XIV Open API is used by storage management applications to communicate with an XIV system.


- [CIM agent overview](#)
A 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.
- [CIM agent components](#)
With a 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.
- [CIM concepts](#)
The common information model (CIM) is an open approach to the management of systems and networks.
- [CIM agent security](#)
The CIM agent can operate in both secure and unsecure modes.

CIM agent overview

A 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 the XIV® system. 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
- Common Information Model (CIM) Specification, Version 2.40.0

- Storage Networking Industry Association (SNIA) Storage Management Initiative Specification (SMI-S) and the Shared Storage Model, 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 storage system through the CIM agent. CIM-compliant applications like the CIM agent can manage the data on your XIV system.

The following sites provide more information about the CIM standards:

- [DMTF Common Information Model \(CIM\) Standards](#)
- [Storage Networking Industry Association Standards](#)

CIM agent components

With a 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 the components of a CIM agent:

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 XIV® system

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.

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.

CIM agent security

The 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.

CIM agent for the XIV system

The CIM agent is preinstalled on the administrative module. The embedded CIM agent is enabled automatically and preconfigured. You can manage the XIV system from the CIM agent that is bundled with the administrative module.

The CIM agent that is embedded on the administrative module has the following limitations:

- The CIM agent can support only the XIV system on which the administrative module is located. The CIM agent is not able to manage any other XIV system.
- The CIM agent must use secure connections over port 5989.
- The CIM agent uses the XIV system user account to authenticate. To manage accounts, you must use the IBM® XIV Storage Management GUI or the IBM XIV command-line interface (XCLI).
- [CIM agent installation overview](#)
The CIM agent is preinstalled on the administrative module. 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.
- [Configuring the CIM agent](#)
The embedded CIM agent does not require configuration changes to manage the XIV system.
- [Enabling the CIM agent](#)
The CIM agent is enabled by default. A watchdog process is monitoring the CIM agent to make sure that it is always running. There are no options to enable, disable, or restart the CIMOM process through the XCLI or IBM XIV Storage Management GUI.

CIM agent installation overview

The CIM agent is preinstalled on the administrative module. 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.

Configuring the CIM agent

The embedded CIM agent does not require configuration changes to manage the XIV® system.

About this task

For the IBM® XIV Storage System version 10.2.4.x, the XIV system user names and passwords are used to authenticate CIM client requests. The XIV system administrator must create user names and passwords that CIM clients can use to attach to the CIM agent. Each user that is created can be in any group. Users that

are created in the *technicians user group* cannot complete any functions. The XIV system administrator can use the XCLI or the IBM XIV Storage Management GUI to create user names and passwords.

Note: CIM uses XML format when communicating with the IBM XIV Storage System, 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:

`& => &`

`< => <`

`> => >`

`' => '`

`" => "`

For the IBM XIV Storage System Gen3 version 11.5 or later, the Multitenancy feature is supported. To support this feature, XIV implements the concept of a domain, which 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 domain will manage different resources.

For the IBM XIV Storage System Gen3 version 11.5 or later, the SMI-S Agent does *not* support the Multitenancy 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 the [IBM\(r\) XIV\(r\) Storage System: Architecture, Implementation, and Usage](#) and the IBM XIV Storage System Product Overview for more information about the architecture of the XIV system.

Enabling the CIM agent

The CIM agent is enabled by default. A watchdog process is monitoring the CIM agent to make sure that it is always running. There are no options to enable, disable, or restart the CIMOM process through the XCLI or IBM® XIV® Storage Management GUI.

XIV Open API component definitions

The XIV Open API elements include schemas, classes, properties, methods, indications, associations, references, and qualifiers.

Elements

The following list describes each type of element:

Schema

A group of classes that are defined to a single namespace. Within the 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

XIV Open API operations are always run within the context of a namespace. A namespace defines the scope over which an XIV Open API schema applies. An XIV 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 XIV 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 XIV Open API implementation that is managed by the XIV 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.

- [CIM agent communication concepts](#)
The CIM agent uses client communication and intrinsic and extrinsic methods of communication.
- [CIM agent communication methods](#)
Client application calls to intrinsic methods can result in 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 CIMOM returns status to the client application.

CIM agent communication concepts

The 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.

CIM agent communication methods

Client application calls to intrinsic methods can result in 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

Functional group	Parameters	Supported or not supported
Association traversal	• Associators • AssociatorNames • References • ReferenceNames	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.

- [GetClass](#)
The GetClass method returns a single class from the target namespace.
- [GetInstance](#)
The GetInstance method returns a single instance from the target namespace.
- [DeleteClass](#)
The DeleteClass method deletes a single class from the target namespace.
- [DeleteInstance](#)
The DeleteInstance method deletes a single instance from the target namespace.
- [CreateClass](#)
The CreateClass method creates class from the target namespace.
- [CreateInstance](#)
The CreateInstance method creates an instance in the target namespace. To use this method, the instance cannot exist.
- [ModifyClass](#)
The ModifyClass method modifies an existing class.
- [ModifyInstance](#)
The ModifyInstance method modifies an existing instance in the target namespace. The instance must exist.
- [EnumerateClasses](#)
The EnumerateClasses method returns a single class from the target namespace.
- [EnumerateClassNames](#)
The EnumerateClassNames method enumerates the names of subclasses of a class that is defined within the target namespace.
- [EnumerateInstances](#)
The EnumerateInstances method enumerates instances of a defined class in the target namespace.
- [EnumerateInstanceNames](#)
The EnumerateInstanceNames method enumerates the names of the instances of a class within a target namespace.
- [ExecuteQuery](#)
The ExecuteQuery method processes a query against the target namespace.
- [Associators](#)
The Associators method enumerates classes or instances that are associated with a specific CIM Object.
- [AssociatorNames](#)
The AssociatorNames method enumerates the names of the classes or instances that are associated with a particular CIM object.
- [References](#)
The References method enumerates the association objects that refer to a particular target class or instance.
- [ReferenceNames](#)
The ReferenceNames method enumerates the association objects that refer to a particular target class or instance.
- [GetProperty](#)
The GetProperty method retrieves a single property value from an instance in the target namespace.
- [SetProperty](#)
The SetProperty method sets a single property value within an instance in the target namespace.
- [GetQualifier](#)
The GetQualifier method retrieves a single qualifier declaration from the target namespace.
- [SetQualifier](#)
The SetQualifier method creates or updates a qualifier declaration in the target namespace.
- [DeleteQualifier](#)
The DeleteQualifier method deletes a single class from the target namespace.
- [EnumerateQualifiers](#)
The EnumerateQualifiers method enumerates qualifier declarations from the target namespace.

GetClass

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

codeblock

coderef

codeph

option

parmname

parml

plentry

pt

pd

synph

syntaxdiagram

groupseq

groupchoice

groupcomp

fragment

fragref

GetInstance

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

DeleteClass

The DeleteClass method deletes a single class from the target namespace.

The DeleteClass method 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.

DeleteInstance

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

CreateClass

The CreateClass method creates class from the target namespace.

The CreateClass method 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.

CreateInstance

The CreateInstance method creates an instance in the target namespace. To use this method, the instance cannot exist.

The CreateInstance method is a standard CIM method. The XIV® 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

ModifyClass

The ModifyClass method modifies an existing class.

The ModifyClass method 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.

ModifyInstance

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

EnumerateClasses

The EnumerateClasses method returns a single class from the target namespace.

Parameters

[Table 1](#) describes the parameters of the EnumerateClasses method.

Table 1. EnumerateClasses method parameters

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

EnumerateClassNames

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

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

EnumerateInstances

The EnumerateInstances method enumerates instances of a defined class in 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

EnumerateInstanceNames

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

ExecuteQuery

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

Associators

The Associators method enumerates classes or instances that are associated with a specific 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

AssociatorNames

The AssociatorNames method enumerates the names of the classes or instances that are associated with a particular 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

References

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

ReferenceNames

The ReferenceNames method enumerates 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

GetProperty

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

SetProperty

The SetProperty method sets a single property value within an instance in the target namespace.

The XIV® 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

GetQualifier

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

SetQualifier

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

DeleteQualifier

The DeleteQualifier method deletes a single class from the target namespace.

The DeleteQualifier method 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.

EnumerateQualifiers

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

Return error codes

The 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_LISTENER_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.
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.

Code	Symbolic Name	Definition
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_UNSUPPORTED_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 diagrams of the CIM agent

The functional diagrams of the CIM agent object model show specific functions that the CIM agent provides. The diagrams also illustrate the architecture of the CIM agent for the XIV® system.

- [Block Server performance profile](#)
Block server performance is the 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 XIV 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 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 SNIA SMI-S specifications define the Mapping and Masking profile to support this function.
- [Indication profile](#)
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 CIM Instances. AlertIndications are used to draw attention of subscribing client applications to the occurrence of an event. Lifecycle Indications are implemented in all XIV CIM Agent releases. Beginning with XIV CIM Agent release 11.0, AlertIndications are implemented.
- [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.
- [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 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 XIV® CIM Agent, statistics are only provided for the volumes and hosts.

[Table 1](#) provides details about the specific metrics that are supported by the XIV arrays and their components. In addition, the table provides a list of ElementType values that correspond to the components of XIV arrays.

Table 1. Block Server Performance metrics

	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

	Equivalent CIM class	Corresponding ElementType value	Properties that are supplied in associated CIM_BlockStorage StatisticalData instances
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.
- [IBMTSDS_PerformanceStatisticsService.GetStatisticsCollection](#)
The IBMTSDS_PerformanceStatisticsService.GetStatisticsCollection method returns a string representation of a set of performance statistics data.
- [Getting volume or host statistics with GetStatisticsCollection](#)
You can use the GetStatisticsCollection method to configure a device.

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 Services Performance SMI-S model for XIV® systems.

Figure 1. Block Services Performance SMI-S model for XIV systems



xiv10228

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 the 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 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 BlockStatisticsManifests 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 XIV 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.

IBMTSDS_PerformanceStatisticsService. GetStatisticsCollection

The IBMTSDS_PerformanceStatisticsService.GetStatisticsCollection method returns a string representation of a set of performance statistics data.

ElementTypes

An array of values that indicates the type of element that returns statistics. Volumes and hosts use the standard value in the 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 BlockStatisticsManifests 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 at the same time, and each ElementType 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 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; }
```

Getting volume or host statistics with GetStatisticsCollection

You can use the GetStatisticsCollection method to configure a device.

About this task

This procedure provides an example recipe for configuring an XIV® device. There are other ways that you can use the XIV Open API to configure XIV devices.

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 ElementType 8) or IBMTSDS_SystemSpecificCollection (or CIM_BlockStatisticsManifest and look for ElementType 4), and then save BulkFormat.
4. Use invokeMethod on GetStatisticsCollection with ElementType 8 for volume or ElementType 4 for host (if ElementType is null, 4 and 8 is assigned to ElementType), StatisticsFormat 2, and the ManifestCollection obtained in the second step, on the PerformanceStatisticsCollection instance that is obtained in the first step.

5. Check the output 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 XIV® 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:

Primordial StoragePool

Logical entities that contains all the available unformatted or unprepared disk capacity on an array. Device Concrete Storage Pools are allocated from the Primordial StoragePool. XIV arrays have only one Primordial StoragePool per array.

Concrete StoragePools

Logical entities that are allocated from the Primordial StoragePool. Concrete StoragePools 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. XIV supports two types of concrete storage pools: VirtualPool and SnapshotPool. VirtualPool is allocated from the Primordial StoragePool directly. It contains volumes and at most one SnapshotPool. SnapshotPool is allocated from VirtualPool and contains snapshots of volumes in the VirtualPool in which this SnapshotPool is located. A VirtualPool without its own SnapshotPool 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 XIV system, 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 XIV Storage System: Architecture, Implementation, and Usage* on the [IBM® Redbooks®](#) website and the *IBM XIV Storage System Product Overview* for more information about the architecture of the XIV system.

- [Block Services object model](#)
The CIM agent is designed to allow for the retrieval of information for each layer of abstraction.
- [IBMTSDS_StorageConfigurationService.CreateOrModifyStoragePool](#)
The CreateOrModifyStoragePool method creates or modifies a VirtualPool or a SnapshotPool. The Primordial StoragePool cannot be created or modified.
- [IBMTSDS_StorageConfigurationService.DeleteStoragePool](#)
DeleteStoragePool requires a single parameter. Specify the reference to the Concrete StoragePool that is to be deleted. Only a single StoragePool can be deleted at a time.
- [IBMTSDS_StorageConfigurationService.CreateOrModifyElementsFromStoragePool](#)
CreateOrModifyElementsFromStoragePool is a vendor-extension method that allows for the creation / modification / moving of multiple volumes in a single method call.
- [IBMTSDS_StorageConfigurationService.CreateOrModifyElementFromStoragePool](#)
CreateOrModifyElementFromStoragePool is the CIM and SMI-S standard method for creating or modifying a StorageVolume from a Concrete StoragePool.
- [IBMTSDS_StorageConfigurationService.ReturnToStoragePool](#)
This method is used to delete a 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.
- [IBMTSDS_StorageConfigurationService.ReturnElementsToStoragePool](#)
This method is used to delete multiple volumes concurrently. You can pass in an array of references of volumes to delete.
- [IBMTSDS_PrimordialStoragePool.GetSupportedSizeRange](#)
CIM_StoragePool has standard methods GetSupportedSizes and GetSupportedSizeRange to find out what sizes to use to create elements from the StoragePool.
- [IBMTSDS_VirtualPool.GetSupportedSizeRange](#)
CIM_StoragePool has standard methods GetSupportedSizes and GetSupportedSizeRange to find out what sizes to use to create elements from the StoragePool.
- [Creating a ConcreteStoragePool](#)
You can use the CreateOrModifyStoragePool method to create a ConcreteStoragePool.
- [Creating a StorageVolume](#)
You can use the CreateOrModifyElementFromStoragePool method to create a StorageVolume.

Block Services object model

The CIM agent is designed to allow for the retrieval of information for each layer of abstraction.

For release 10.2.4.x, SMI-S specification 1.4 is fully supported. .

For the 11.6 release, 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 XIV® systems.

Figure 1. Block Services SMI-S model for XIV systems



A StorageSetting is a class that contains properties to specify the quality of service, such as DataRedundancy or parity layout. When you create a Storage Volume or a Concrete StoragePool, a StorageSetting is supplied as the Goal parameter for the appropriate method.

XIV CIM Agent has three CIM_StorageSetting instances defined:

- IBMTSDS_VirtualPoolSetting.InstanceID="IBMTSDS:IBM XIV Virtual Storage Pool Setting": This StorageSetting is supplied as the Goal parameter for the CreateOrModifyStoragePool method to create or modify VirtualPool.
- IBMTSDS_SnapshotPoolSetting.InstanceID="IBMTSDS:IBM XIV Snapshot Pool Setting": This StorageSetting is supplied as the Goal parameter for the CreateOrModifyStoragePool method to create or modify SnapshotPool.
- IBMTSDS_DataTypeSetting.InstanceID="IBMTSDS:XIVBlockSize": This StorageSetting is supplied as the Goal parameter for the CreateOrModifyElementFromStoragePool and CreateOrModifyElementsFromStoragePool method to create or modify Storage Volumes with size in terms of GB. The system allocates the soft volume size as the minimum number of discrete 17 GB increments needed to meet the requested volume size.
- IBMTSDS_DataTypeSetting.InstanceID="IBMTSDS:SystemBlockSize": This StorageSetting is supplied as the Goal parameter for the CreateOrModifyElementFromStoragePool and CreateOrModifyElementsFromStoragePool method to create or modify Storage Volumes 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 Storage Pool as the minimum number of discrete 17 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 a CIM_StorageConfigurationService and the result class CIM_StorageConfigurationCapabilities. The SupportedSynchronousActions and SupportedAsynchronousActions properties identifies whether the action is initiated synchronously or asynchronously. XIV CIM Agent supports the following synchronous actions in IBMTSDS_StorageConfigurationService class:

Table 1. Synchronous actions

Synchronous actions	XIV CIM functions
StoragePool Creation	CreateOrModifyStoragePool
StoragePool Modification	CreateOrModifyStoragePool
StoragePool Deletion	DeleteStoragePool

Synchronous actions	XIV CIM functions
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 StorageVolumes can be created:

IBMTSDS_PrimordialStoragePool.GetSupportedSizeRange

Determines the supported physical size of Concrete StoragePools that can be created in Primordial StoragePool.

IBMTSDS_VirtualPool.GetSupportedSizeRange

Determines the supported soft size of StorageVolumes that can be created in a Concrete StoragePool.

IBMTSDS_StorageConfigurationService. CreateOrModifyStoragePool

The CreateOrModifyStoragePool method creates or modifies a VirtualPool or a SnapshotPool. The Primordial StoragePool cannot be created or modified.

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, only the object reference of IBMTSDS_PrimordialStoragePool is allowed. All VirtualPools are created in the Primordial StoragePool. To create a SnapshotPool, only the object reference of IBMTSDS_VirtualPool is allowed because SnapshotPool is created in a VirtualPool.

Note: The 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 represents the StorageSetting of the pool to be created. To create a VirtualPool, only the object reference IBMTSDS_VirtualPoolSetting is allowed; to create a SnapshotPool, only the object reference IBMTSDS_SnapshotPoolSetting is allowed. The Goal parameter also specifies the requested hard size 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, -, _, . and ~ characters. And 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 soft size of the pool. **Null** or **0** is not allowed for pool creation. **0** for pool modification means that the pool's size is not changed. As an output parameter, Size specifies the soft size achieved. If it is not possible to create a pool of the wanted size, a size is returned with size value set to the nearest supported size.

IBMTSDS_StorageConfigurationService. DeleteStoragePool

DeleteStoragePool requires a single parameter. Specify the reference to the Concrete StoragePool that is to be deleted. Only a single StoragePool can be deleted at a time.

You must delete all volumes in the current pool before you delete a pool. The Primordial StoragePool cannot be deleted.

IBMTSDS_StorageConfigurationService. CreateOrModifyElementsFromStoragePool

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.

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 StorageVolumes to be created. This parameter must be null when you modify StorageVolumes. As output, represents the number of volumes that are created if successful. Or creation was not successful because too many volumes and not enough capacity on the VirtualPool, represents the number of volumes that can be created.

ElementNames

An array of the ElementNames to assign to the various volumes that are being created or modified. Volumes 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 Quantity when you create volumes. If not **NULL**, the length must be the same as the length of TheElements when you modify volumes. For example, to create two volumes 'testVolume_3' and 'testVolume_4', set the parameter's value as ['testVolume_3' , 'testVolume_4'] and leave the value of FirstSuffix 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 this parameter's value as ['testVolume'] and set value of FirstSuffix parameter as 3.

FirstSuffix

A uint16 that 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, it must be null. Note: The FirstSuffix parameter is not in the CIM or 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.

IBMTSDS_StorageConfigurationService. CreateOrModifyElementFromStoragePool

CreateOrModifyElementFromStoragePool is the CIM and SMI-S standard method for creating or modifying a StorageVolume from a Concrete StoragePool.

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.

The following definitions summarize how the method 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 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 representing the DataType 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.

IBMTSDS_StorageConfigurationService. ReturnToStoragePool

This method is used to delete a 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.

IBMTSDS_StorageConfigurationService. ReturnElementsToStoragePool

This 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 CIM or SMI-S schema. This method is an IBM® extension.

IBMTSDS_PrimordialStoragePool.GetSupportedSizeRange

CIM_StoragePool has standard methods GetSupportedSizes and GetSupportedSizeRange to find out what sizes to use to create elements from the StoragePool.

For IBMTSDS_PrimordialStoragePool, only GetSupportedSizeRange is supported. It can be used to retrieve the supported physical size of 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 storage pool is zero.

IBMTSDS_VirtualPool.GetSupportedSizeRange

CIM_StoragePool has standard methods GetSupportedSizes and GetSupportedSizeRange to find out what sizes to use to create elements from the StoragePool.

This method can be used to retrieve the supported physical 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 storage pool is zero.

Creating a ConcreteStoragePool

You can use the CreateOrModifyStoragePool method to create a ConcreteStoragePool.

About this task

This procedure provides an example recipe for configuring an XIV® device. There are other ways that you can use the XIV Open API to configure XIV devices.

Procedure

1. Call IBMTSDS_PrimordialStoragePool.GetSupportedSizeRange with ElementType set to Storage Pool. 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 that was created.

Creating a StorageVolume

You can use the CreateOrModifyElementFromStoragePool method to create a StorageVolume.

About this task

This procedure provides an example recipe for configuring an XIV® device. There are other ways that you can use the XIV Open API to configure XIV devices.

Procedure

1. Use associatorNames from the Concrete StoragePool that was created to CIM_StorageCapabilities to get matching instance of IBMTSDS_VirtualPoolCapabilities.
2. Use associatorNames from IBMTSDS_VirtualPoolCapabilities 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 IBMTSDS_VirtualPool created in Step 1, IBMTSDS_DataTypeSetting, Size, and ElementName 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 CIM objects.

A client can obtain information about iSCSI capabilities and settings of an XIV® 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 XIV SMI-S Provider.

- **[iSCSI Target Ports Object Model](#)**
The CIM agent is designed to allow for the retrieval of information for each layer of abstraction.
- **[IBMTSDS_iSCSIConfigurationService.CreateIPProtocolEndpoint](#)**
The IBMTSDS_iSCSIConfigurationService.CreateIPProtocolEndpoint method creates an IP endpoint.
- **[IBMTSDS_iSCSIConfigurationService.ModifyIPProtocolEndpoint](#)**
The IBMTSDS_iSCSIConfigurationService.ModifyIPProtocolEndpoint method modifies an IP endpoint.
- **[IBMTSDS_iSCSIConfigurationService.DeleteIPProtocolEndpoint](#)**
This method is used to delete an IP PortocolEndPoint. The only input parameter that is required is a reference to the PortocolEndPoint that you want to delete. Only one PortocolEndPoint can be deleted at a time.
- **[Creating an IP endpoint](#)**
You can use CreateIPProtocolEndpoint method to create an IP endpoint.
- **[Modifying an IP endpoint](#)**
You can use ModifyIPProtocolEndpoint method to create an IP endpoint.
- **[Deleting an IP endpoint](#)**
You can use the DeleteIPProtocolEndpoint method to delete an IP endpoint.

iSCSI Target Ports Object Model

The CIM agent is designed to allow for the retrieval of information for each layer of abstraction.

See [Figure 1](#) for the iSCSI Target Ports SMI-S model for XIV® 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 XIV 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.
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.

iSCSI term	CIM class name	Description
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 is considered a separate Network Entity, which instance associated to an instance of CIM_iSCSICapabilities through the CIM_ElementCapabilities association 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

The iSCSI Node is represented by instances of CIM_SCISIProtocolController. This class can have many subclasses. The Name and NameFormat properties can be used to determine whether a CIM_SCISIProtocolController is an iSCSI Node. All CIM_SCISIProtocolController instances that represent iSCSI Nodes are required to have a Name property whose value is an iSCSI Name. Therefore, the NameFormat property must have a value of 3 (which maps to iSCSI Name). On the XIV arrays, the iSCSI Node is represented by instances of IBMTSDS_iSCSIProtocolController. It is a subclass of CIM_SCISIProtocolController, which is associated to the first-level system (IBMTSDS_StorageSystem) through CIM_SystemDevice. For XIV arrays, each IBMTSDS_StorageSystem has one Node for each IP Interface. The iSCSI Names 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 on a XIV 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 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 to the CIM_iSCSIProtocolEndpoint instances.

Network Portal

There is a single network portal per IP Interface on a XIV 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:

- CIM_IPProtocolEndpoint contains the IP address of the Network Portal in the IPv4Address property. The IPv6Address property is not currently supported.
- CIM_TCPProtocolEndpoint instance contains the TCP Port of the Network Portal in the PortNumber property.
- CIM_EthernetPort instance contains the physical properties of the Network Portal in the PermanentAddress, DeviceID, FullDuplex, Speed, and OtherIdentifyingInfo properties.

IBMTSDS_iSCSIConfigurationService. CreateIPProtocolEndpoint

The IBMTSDS_iSCSIConfigurationService.CreateIPProtocolEndpoint method creates an IP endpoint.

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.

IBMTSDS_iSCSIConfigurationService. ModifyIPProtocolEndpoint

The IBMTSDS_iSCSIConfigurationService.ModifyIPProtocolEndpoint method modifies an IP endpoint.

The method ModifyIPProtocolEndpoint is not in the CIM/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.

IBMTSDS_iSCSIConfigurationService. DeleteIPProtocolEndpoint

This method is used to delete an IP PortocolEndPoint. The only input parameter that is required is a reference to the PortocolEndPoint that you want to delete. Only one PortocolEndPoint can be deleted at a time.

The method DeleteIPProtocolEndpoint is not in the CIM/SMI-S schema but an IBM® extension.

Creating an IP endpoint

You can use CreateIPProtocolEndpoint method to create an IP endpoint.

About this task

This procedure provides an example recipe for configuring an XIV® device. There are other ways that you can use the XIV Open API to configure XIV devices.

Procedure

1. Obtain the reference (CIMObjectPath) of an IBMTSDS_iSCSIConfigurationService instance. The instance is associated with the IBMTSDS_StorageSystem 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.

Modifying an IP endpoint

You can use ModifyIPProtocolEndpoint method to create an IP endpoint.

About this task

This procedure provides an example recipe for configuring an XIV® device. There are other ways that you can use the XIV Open API to configure XIV devices.

Procedure

1. Obtain the reference (CIMObjectPath) of an IBMTSDS_iSCSIConfigurationService instance that is associated with the "IBMTSDS_StorageSystem" 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.

Deleting an IP endpoint

You can use the DeleteIPProtocolEndpoint method to delete an IP endpoint.

About this task

This procedure provides an example recipe for configuring an XIV® device. There are other ways that you can use the XIV Open API to configure XIV devices.

Procedure

1. Obtain the reference (CIMObjectPath) of an IBMTSDS_iSCSIConfigurationService instance. Ensure that the reference is associated with the IBMTSDS_StorageSystem 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.

Masking and Mapping profile

You can control which client hosts can see and access storage volumes by mapping and masking client hosts. The SNIA 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 the XIV® system 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 XIV system architecture for masking and mapping involves the following objects on two kinds of devices:

- Objects on the XIV system
 - 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 the XIV system uses masking and mapping.

Figure 1. Masking and mapping physical model in the XIV system



The following list defines each of the objects that are shown in [Figure 1](#).

Storage device I/O ports

Ports on the XIV system 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 the XIV system 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® XIV Storage Management GUI, XCLI, or CIM agent.

Clusters

Groups of hosts. For the XIV system, 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 XIV Storage Management GUI, XCLI, or CIM agent. Clusters are required when more than one host can access the same volume.

After you set up the physical network and created a host object, add Fibre Channel ports or iSCSI ports to the host. Map the volumes to the host so the volume can be used for I/O operations.

Map 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.

Map 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 XIV Storage System: Architecture, Implementation, and Usage* and the *IBM XIV Storage System Product Overview* for more information about the architecture of the XIV system.


- [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.
- [IBMTSDS_StorageHardwareIDManagement.CreateStorageHardwareID](#)
You can use the CreateStorageHardwareID method to create a StorageHardwareID. This method creates a host in the XIV device that contains the initiator port that is specified by StorageID.
- [IBMTSDS_StorageHardwareIDManagement.DeleteStorageHardwareID](#)
Use the DeleteStorageHardwareID method to delete StorageHardwareIDs.
- [IBMTSDS_StorageHardwareIDManagement.CreateHardwareIDCollection](#)
You can use the CreateHardwareIDCollection method to create HardwareIDCollections.
- [IBMTSDS_StorageHardwareIDManagement.AddHardwareIDsToCollection](#)
You can use the AddHardwareIDsToCollection method to add StorageHardwareIDs (representing initiator ports) into SystemSpecificCollection (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.
- [IBMTSDS_ControllerConfigurationService.ExposePaths](#)
Use the ExposePaths method to assign volumes to a StorageHardwareID representing a host or cluster in XIV.
- [IBMTSDS_ControllerConfigurationService.HidePaths](#)
The HidePaths method is the inverse of ExposePaths, in other words, removing access to volumes from a StorageHardwareID.
- [IBMTSDS_ControllerConfigurationService.DeleteProtocolController](#)
The DeleteProtocolController method deletes an existing ProtocolController. In XIV, the corresponding host and all the host ports that are contained by the host are deleted.
- [Sample recipe](#)
The following task is a sample recipe for configuring an XIV device. It is only a sample; there are many ways to use the XIV Open API to configure an XIV device.

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 XIV® 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 CIM agent installation directory.

Table 1. Masking and mapping classes

CIM class name	Description	XIV CIM class name
CIM_FCPort	A Fibre Channel port on the XIV system.	IBMTSDS_FCPort
CIM_EthernetPort	An Ethernet port on the XIV 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 XIV system	IBMTSDS_SEVolume

CIM class name	Description	XIV CIM class name
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

XIV CIM Agent release 10.2.4 provides a full implementation of masking and mapping subprofile including the following items:

- iSCSI port support
- StorageHardwareID creation and deletion
- HardwareIDCollection creation and modification
- Expose paths and hide paths

iSCSI Port support

By default, the XIV Storage System ships equipped with six iSCSI ports. Three of the interface modules support iSCSI, with two ports in each module. You can refer to XIV Redbook for detailed information. XIV CIM Agent 10.2.4 adds the full iSCSI support by implementation of the iSCSI target ports profile and the masking and mapping profile.

StorageHardwareID manipulation

A StorageHardwareID represents the host-side initiator that logs in to a storage area network. The XIV SMI-S provider supports host-side initiators that contain Fibre Channel or iSCSI host ports. A StorageHardwareIDManagementService allows a CIM client to locate StorageHardwareIDs known to a storage array, and create, modify, and delete StorageHardwareIDs.

- CreateStorageHardwareID() method creates StorageHardwareID. On XIV, a host is created which contains the specified initiator port.
- DeleteStorageHardwareID() deletes a StorageHardwareID. On XIV, the initiator port is removed from the host that contains it.

HardwareIDCollection manipulation

A HardwareIDCollection represents a collection of StorageHardwareID (host-side initiators) that logs in to a storage area network. A StorageHardwareIDManagementService allows a CIM client to locate HardwareIDCollections known to a storage array, and create HardwareIDCollection and remove HardwareIDs from HardwareIDCollection.

- CreateHardwareIDCollection method creates a HardwareIDCollection with specified ElementName and Setting.
- AddHardwareIDsToCollection adds StorageHardwareIDs into a HardwareIDCollection.

Note: The methods ModifyHardwareIDCollection, DeleteHardwareIDCollection, and RemoveHardwareIDsFromCollection are not in the CIM or SMI-S schema. These methods are an IBM® extension.

Expose paths and hide paths

A SCSIProtocolController represents a view of volumes that can be assigned access to a StorageHardwareID. You can use a ControllerConfigurationService to locate SCSIProtocolControllers as well as ExposePaths and HidePaths from StorageHardwareIDs to SCSIProtocolControllers.

A privilege determines the access rights between the StorageHardwareID and the SCSIProtocolController. The XIV SMI-S Provider supports read/write access. A privilege can be located by using the PrivilegeManagementService.

IBMTSDS_StorageHardwareIDManagement. CreateStorageHardwareID

You can use the CreateStorageHardwareID method to create a StorageHardwareID. This method creates a host in the XIV device that contains the initiator port that is specified by StorageID.

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 Standard, HPUX, and z/VM are supported by XIV CIM Agent. If it is null, Standard is used by default.

IDType

An input parameter that represents the type of port. To create a HardwareID representing a Fibre Channel port, specify 2; to create a HardwareID representing an iSCSI port, specify 5.

StorageID

An input parameter that represents the ID of the port. If IDType is 2, specify it as the port's WWN; if IDType is 5, specify it as the port's IQN.

HardwareID

An output parameter; the reference to the created StorageHardwareID.

IBMTSDS_StorageHardwareIDManagement. DeleteStorageHardwareID

Use the DeleteStorageHardwareID method to delete StorageHardwareIDs.

HardwareID

An input parameter that represents a reference to the StorageHardwareID to delete.

IBMTSDS_StorageHardwareIDManagement. CreateHardwareIDCollection

You can use the CreateHardwareIDCollection method to create HardwareIDCollections.

ElementName *(required)*

An input parameter that represents a meaningful name for the HardwareIDCollection 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 HardwareIDCollections in the system. If not specified, CIM generates a random name.

HardwareIDs

An input parameter that represents the ID of the ports that are contained by the HardwareIDCollection. 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 is created.

HardwareIDCollection

An output parameter that represents the reference to the created HardwareIDCollection.

IBMTSDS_StorageHardwareIDManagement. AddHardwareIDsToCollection

You can use the AddHardwareIDsToCollection method to add StorageHardwareIDs (representing initiator ports) into SystemSpecificCollection (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.

HardwareIDs *(required)*

An input parameter that represents an array of strings that contain StorageIDs of StorageHardwareID instances that become members of the HardwareIDCollection.

HardwareIDCollection *(required)*

An input parameter that represents the reference to the SystemSpecificCollection into which the HardwareIDs are added.

IBMTSDS_ControllerConfigurationService. ExposePaths

Use the ExposePaths method to assign volumes to a StorageHardwareID representing a host or cluster in XIV®.

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 SCSIProtocolControllers (SPCs). It supports creating or modifying SPCs depending on different specified parameters.

Table 1. ExposePaths parameters and use cases

Parameters/ Use cases	LUNames	Initiator PortIDs	TargetPortIDs	DeviceNumbers	DeviceAccesses	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

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), which is the WWN of an XIV 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.

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. If not specified, the CIM Agent assigns unused LUN IDs for the operation.

DeviceAccesses

A unit16 array input parameter that represents the access rights to give to the StorageHardwareID to the volumes specified in LUNames. Read/write access is supported. It must be an array the same size as 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 to be modified. As an output parameter, it represents the SCSIProtocolController that is being modified or created.

IBMTSDS_ControllerConfigurationService. HidePaths

The HidePaths method is the inverse of ExposePaths, in other words, removing access to volumes from a StorageHardwareID.

Table 1. HidePaths parameters and use cases

Parameters/ Use cases	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

LUNames (*required for removing access to volumes*)

A string array input parameter that represents the volumes to unmap from the StorageHardwareID. They are not references to volume instances, but are strings that match the name property of the IBMTSDS_SEVolume (CIM_StorageVolume).

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 to be modified. As an output parameter, it represents the SCSIProtocolController that was modified or deleted.

IBMTSDS_ControllerConfigurationService.DeleteProtocolController

The DeleteProtocolController method deletes an existing ProtocolController. In XIV, the corresponding host and all the host ports that are contained by the host are deleted.

ProtocolController (*required*)

An input parameter that represents the ProtocolController to be deleted.

DeleteUnits

An input parameter that represents the requirement to delete the volumes that are mapped to the ProtocolController. If true is specified, the mapped volumes are deleted.

Sample recipe

The following task is a sample recipe for configuring an XIV device. It is only a sample; there are many ways to use the XIV Open API to configure an XIV device.

1. Map FC and iSCSI ports to volumes:
 - a. Attain WWN of host FC ports, IQN of host iSCSI ports, and volume names.
 - b. 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.
2. Add volumes to an existing SCSIProtocolController:
 - a. Attain the host reference and volume names.
 - b. Find the IBMTSDS_SCSIProtocolController instance that matches the selected host.
 - c. Use InvokeMethod on IBMTSDS_ControllerConfigurationService.ExposePaths with LUNames set to volume names, DeviceAccesses set to 2, and ProtocolControllers set to the SCSIProtocolController.
3. Add initiator ports to an existing SPC:
 - a. Attain IDs of initiator ports.
 - b. Find the IBMTSDS_SCSIProtocolController instance.
 - c. Use InvokeMethod on IBMTSDS_ControllerConfigurationService.ExposePaths with InitiatorPortIDs set to the initiator port IDs, DeviceAccesses set to 2, and ProtocolControllers set to the SCSIProtocolController.
4. Remove volumes from an SPC:
 - a. Attain the SPC reference and volume names.
 - b. Find the IBMTSDS_SCSIProtocolController instance that matches the selected host.
 - c. Use InvokeMethod on IBMTSDS_ControllerConfigurationService.HidePaths with LUNames set to volume names, and ProtocolControllers set to the SCSIProtocolController.
5. Remove initiator ports from an SPC:
 - a. Attain the SPC reference and IDs of the initiator port.
 - b. Find the IBMTSDS_SCSIProtocolController instance.
 - c. Use InvokeMethod on IBMTSDS_ControllerConfigurationService.HidePaths with InitiatorPortIDs set to the IDs of the initiator ports, and ProtocolControllers set to the SCSIProtocolController.

Indication profile

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 CIM Instances. AlertIndications are used to draw attention of subscribing client applications to the occurrence of an event. Lifecycle Indications are implemented in all XIV® CIM Agent releases. Beginning with XIV CIM Agent release 11.0, AlertIndications are implemented.

Supported lifecycle indications by class

[Table 1](#) identifies the class of objects that can be monitored by using the specified indication type for XIV arrays.

Table 1. Indication types and object classes

Class name	Supported lifecycle indications
IBMTSDS_SEVolume	IBMTSDS_InstCreation
	IBMTSDS_InstDeletion
	IBMTSDS_InstModification

Class name	Supported lifecycle indications
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 XIV CIM Agent configuration function is successfully completed. If SMI-S client starts the XIV 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, XIV CIM Agent periodically fetches XIV events and sends related lifecycle indications. If a StorageVolume is created on an XIV array through the XIV GUI, an SMI-S client receives a related InstCreation indication with the new IBMTSDS_SEVolume instance from the provider. By default, the XIV event fetch interval is 30 seconds.

Alert indications

The XIV CIM Agent periodically fetches XIV 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 XIV events fetching interval is 30 seconds.

Note: The XIV 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 acknowledgement that a write to the source element was processed by the target element. Until it receives this acknowledgement, it does not accept another I/O from the host. In asynchronous mode, the writer does not wait for this acknowledgement 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.
- [**Group management**](#)
The CIM agent supports group manipulation, including consistency group and snapshot group.
- [**IBMTSDS_ReplicationService.CreateGroup**](#)
The CreateGroup method creates a consistency group.
- [**IBMTSDS_ReplicationService.DeleteGroup**](#)
The DeleteGroup method deletes an existing consistency group.
- [**IBMTSDS_ReplicationService.AddMembers**](#)
The AddMembers method adds volumes to a consistency group.
- [**IBMTSDS_ReplicationService.RemoveMembers**](#)
The RemoveMembers method removes volumes from a consistency group.
- [**Replication management**](#)
The CIM agent version 11.1 and later supports local replication manipulation, including local snapshot and clone. The CIM agent version 11.6 and later also supports remote replication manipulation, including volume and consistency group mirroring.
- [**IBMTSDS_ReplicationService.CreateElementReplica**](#)
The CreateElementReplica method creates a snapshot of a volume, duplicate snapshot of a snapshot, clone of a volume, and clone of a snapshot.
- [**IBMTSDS_ReplicationService.CreateGroupReplica**](#)
The CreateGroupReplica method creates a snapshot group for a consistency group.
- [**IBMTSDS_ReplicationService.ModifyReplicaSynchronization**](#)
For local replication service, use the 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 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).
- [**IBMTSDS_ReplicationService.ModifyListSynchronization**](#)
This method is a batch operation of the IBMTSDS_ReplicationService:ModifyReplicaSynchronization method, which modifies a list of synchronizations.
- [**IBMTSDS_ReplicationService.GetAvailableTargetElements**](#)
The GetAvailableTargetElements method retrieves all of the candidate target elements for the supplied source element.
- [**IBMTSDS_ReplicationService.GetReplicationRelationships**](#)
The GetReplicationRelationships method retrieves all of the snapshot or mirror relationships on a XIV device.
- [**IBMTSDS_ReplicationService.GetReplicationRelationshipInstances**](#)
This method is similar to GetReplicationRelationships, except that it returns strings of references of snapshots, snapshot groups, or mirror relationships.
- [**IBMTSDS_ReplicationService.GetPeerSystems**](#)
Use this method to get (or start a job to get) 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 XIV Storage Systems.
- [**IBMTSDS_ReplicationService.CreateGroupReplicaFromElements**](#)
Use this method to create (or start a job to create) new storage objects that are replicas of the specified source storage objects (SourceElements).

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 XIV® systems





Figure 2. Replication Services (Remote) SMI-S model for XIV systems



Replication Service capabilities

Use XIV terminology to map to SMI-specific terminology:

Table 1. Mapping XIV terminology to SMI terminology

XIV Terminology	SMI Terminology	Description	XIV 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 11.1 or later

XIV Terminology	SMI Terminology	Description	XIV CIM Agent support
Volume copy	Synchronous Clone Local	A full, "Point-In-Time", unassociated local copy of the source element.	version 11.1 or later
Synchronous Remote Mirroring	Synchronous Mirror Remote	A synchronized remote copy of the source element.	version 11.6 or later
Asynchronous Remote Mirroring	Asynchronous Mirror Remote	An asynchronous remote copy of the source element.	version 11.6 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.

Group management

The CIM agent supports group manipulation, including consistency group and snapshot group.

Table 1. Group management classes

XIV	SMI class	XIV 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

See [Figure 1](#).

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

IBMTSDS_ReplicationService. CreateGroup

The CreateGroup method creates a consistency group.

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.

IBMTSDS_ReplicationService. DeleteGroup

The DeleteGroup method deletes an existing consistency group.

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.

IBMTSDS_ReplicationService. AddMembers

The AddMembers method adds volumes to a consistency group.

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.

IBMTSDS_ReplicationService. RemoveMembers

The RemoveMembers method removes volumes from a consistency group.

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

The CIM agent version 11.1 and later supports local replication manipulation, including local snapshot and clone. The CIM agent version 11.6 and later 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

XIV	SMI class	XIV 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 Primary XIV and Secondary XIV respectively; see [Figure 2](#). At the same time, the IBMTSDS_SourceConsistencyGroupSynchronized and IBMTSDS_TargetConsistencyGroupSynchronized 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.

IBMTSDS_ReplicationService. CreateElementReplica

The CreateElementReplica method creates a snapshot of a volume, duplicate snapshot of a snapshot, clone of a volume, and clone of a snapshot.

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, CIM selects the first available pool on the device to create the clone.

IBMTSDS_ReplicationService. CreateGroupReplica

The CreateGroupReplica method creates a snapshot group for a consistency group.

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.

IBMTSDS_ReplicationService. ModifyReplicaSynchronization

For local replication service, use the 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 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).

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).

IBMTSDS_ReplicationService. ModifyListSynchronization

This method is a batch operation of the IBMTSDS_ReplicationService:ModifyReplicaSynchronization method, which modifies a list of synchronizations.

IBMTSDS_ReplicationService. GetAvailableTargetElements

The `GetAvailableTargetElements` method retrieves all of the candidate target elements for the supplied source element.

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.

IBMTSDS_ReplicationService. GetReplicationRelationships

The `GetReplicationRelationships` method retrieves all of the snapshot or mirror relationships on a XIV device.

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.

IBMTSDS_ReplicationService. GetReplicationRelationshipInstances

This method is similar to `GetReplicationRelationships`, except that it returns strings of references of snapshots, snapshot groups, or mirror relationships.

IBMTSDS_ReplicationService. GetPeerSystems

Use this method to get (or start a job to get) 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 XIV Storage Systems.

IBMTSDS_ReplicationService. CreateGroupReplicaFromElements

Use this 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 `CreateGroup` and `CreateGroupReplica`, in one call. This method creates mirrors for volumes and consistency groups and then adds the mirrored volumes into the mirrored consistency groups.

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. If TargetGroup and TargetPool are specified at the same time, and TargetGroup is not in the TargetPool, TargetGroup will be moved to TargetPool.

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.

Table 1. SMI Indications

SMI 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 XIV 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, 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 XIV® systems



For IBMTSDS_MirrorJob, if all the operations completed successfully, the JobState will be 7 (Completed), and the OperationStatus will be 2 (Ok), 17 (Completed). If not, the JobState will be 10 (Exception), and the OperationStatus 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 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 the IBM XIV Storage System, thin provisioning is a pool definition (not a volume definition). It defines the pool as a thin provisioned pool and then assigns different soft and hard sizes to it. From the volume perspective, it is the same. You define the volume size (no soft size or hard size distinction), and it takes the size from the soft size of the pool. Upon creation of a volume in a regular or thin provisioned pool, no pre-allocation of hard capacity is done. Only soft capacity is taken from the soft 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.
- [IBMTSDS_PrimalStoragePoolCapabilities.CreateSetting](#)
You can use the CreateSetting method to create an "IBMTSDS_VirtualPoolSetting" instance.
- [IBMTSDS_VirtualPoolSetting.modify](#)
You can use the modify method to modify the properties of a newly created IBMTSDS_VirtualPoolSetting instance.
- [IBMTSDS_VirtualPoolSetting.delete](#)
You can use the delete method to delete the newly created IBMTSDS_VirtualPoolSetting instance. The default instance of IBMTSDS_VirtualPoolSetting cannot be deleted.
- [IBMTSDS_StorageConfigurationService.CreateOrModifyStoragePool](#)
You can use the CreateOrModifyStoragePool method to create or modify a regular VirtualPool or a thin provisioning VirtualPool.
- [Creating a thin provisioning pool](#)
You can use the CreateOrModifyStoragePool method to create a thin provisioning pool.
- [Thin provisioning indications](#)
The XIV 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 SMI-S model for XIV® systems.

Figure 1. Thin Provisioning SMI-S model for XIV systems



As with the Block Services profile, the CIM agent has three 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 is supplied as the Goal parameter for the CreateOrModifyStoragePool method to create or modify a regular VirtualPool.

IBMTSDS_VirtualPoolSetting.InstanceID="IBMTSDS:IBM.2810-1310129-112714143137626"

This StorageSetting is supplied as the Goal parameter for the CreateOrModifyStoragePool method to create or modify a thin provisioning VirtualPool.

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 regular VirtualPool or thin provisioning VirtualPool

IBMTSDS_PrimordialStoragePoolCapabilities. CreateSetting

You can use the CreateSetting method to create an "IBMTSDS_VirtualPoolSetting" instance.

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.

IBMTSDS_VirtualPoolSetting. modify

You can use the modify method to modify the properties of a newly created IBMTSDS_VirtualPoolSetting instance.

ThinProvisionedInitialReserve *(required if ThinProvisionedPoolType is null)*

As an input, represents the hard size of the newly created pool to be used as the Goal parameter for the CreateOrModifyStoragePool method.

ThinProvisionedPoolType *(required if ThinProvisionedInitialReserve is null)*

As an input, represents the pool type that you want to create to use as the goal parameter for the CreateOrModifyStoragePool method. 7 indicates that the newly created pool is a regular pool, and 8 indicates that the newly created pool is a thin provisioning pool.

Cop

As an input, represents the CIM Object Paths (COPs) of the newly created IBMTSDS_VirtualPoolSetting instance to modify.

IBMTSDS_VirtualPoolSetting. delete

You can use the delete method to delete the newly created IBMTSDS_VirtualPoolSetting instance. The default instance of IBMTSDS_VirtualPoolSetting cannot be deleted.

Cop

As an input, represents the CIM Object Paths (COPs) of the newly created IBMTSDS_VirtualPoolSetting instance to delete.

IBMTSDS_StorageConfigurationService. CreateOrModifyStoragePool

You can use the CreateOrModifyStoragePool method to create or modify a regular VirtualPool or a thin provisioning VirtualPool.

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

InPools specifies from which pool to create the pool. To create a VirtualPool, only the object reference of IBMTSDS_PrimordialStoragePool is allowed for this parameter since all VirtualPools are created in the Primordial StoragePool. To create a SnapshotPool, only the object reference of IBMTSDS_VirtualPool is allowed for this parameter if specified since SnapshotPool is created in a VirtualPool.

Note: The 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 represents the StorageSetting of the pool to be created. To create a VirtualPool, only the object reference of IBMTSDS_VirtualPoolSetting is allowed. To create a SnapshotPool, only the object reference of IBMTSDS_SnapshotPoolSetting is allowed. The Goal parameter also specifies the requested hard size 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 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 soft size of the Pool. Null or 0 is not allowed for pool creation. 0 for pool modification indicates that the pool's size is not changed. As an output parameter, Size specifies the soft size achieved. If the value entered is not supported, the size value is set to the nearest supported size.

Creating a thin provisioning pool

You can use the CreateOrModifyStoragePool method to create a thin provisioning pool.

About this task

This procedure provides an example for configuring an XIV® device. There are other ways that you can use the XIV Open API to configure XIV devices.

Procedure

1. Call IBMTSDS_PrimordialStoragePool.GetSupportedSizeRange with ElementType set to **Storage Pool**. If the return code is 0 (success), specify a size.
2. Call invokeMethod on IBMTSDS_PrimordialStoragePoolCapabilities.CreateSetting and specify the SettingType.
3. Call enumerateInstances on IBMTSDS_VirtualPoolSetting, and then save the CIMObjectPath of this newly created instance.
4. Call invokeMethod on IBMTSDS_VirtualPoolSetting.modify and specify the IBMTSDS_VirtualPoolSetting obtained in step 3 to Cop, pool hard size to ThinProvisionedInitialReserve, and pool type to ThinProvisionedPoolType.
Note: For the thin provisioning pool, the hard size must be less than the soft size.
5. Call invokeMethod on IBMTSDS_StorageConfigurationService.CreateOrModifyStoragePool and specify the IBMTSDS_VirtualPoolSetting obtained in step 3 to Goal and ElementName to Size.
6. Save the pool output parameter. It is the reference to the thin provisioning VirtualPool that was created.

Thin provisioning indications

The XIV® CIM Agent supports three thin provisioning profile indications.

Table 1. Thin provisioning profile Indications

CIM Indication	XIV Event	Description
CAPACITY_WARNING	STORAGE_POOL_VOLUME_USAGE_TOO_HIGH	Thin provisioned <i>Volume or Pool</i> with the identifier <i>Volume or Pool ID</i> capacity in use near available limit.
CAPACITY_CRITICAL	STORAGE_POOL_EXHAUSTED	Thin provisioned <i>Volume or Pool</i> with the identifier <i>Volume or Pool ID</i> capacity in use exceeded available limit.
CAPACITY_OKAY	STORAGE_POOL_VOLUME_USAGE_BACK_TO_NORMAL	Thin provisioned <i>Volume or Pool</i> with the identifier <i>Volume or Pool ID</i> capacity condition cleared.

Conformance tests

All SNIA official Certification Test Programs (CTP) and Microsoft System Center Virtual Machine Manager (SCVMM) Storage Automation tests are complete.

SNIA official CTP tests

Table 1. SNIA CTP results

XIV CIM Agent release	Official CTP Test Suite	Official CTP Test Results
10.2.4	• SMI-S version 1.4 • Test version 1.4.0.24	http://www.snia.org/ctp/conformingproviders/ibm.html#sftw3
11.0	• SMI-S version 1.4 • Test version 1.4.0.24	http://www.snia.org/ctp/conformingproviders/ibm.html#sftw3
11.1	• SMI-S version 1.4 • Test version 1.4.0.24	http://www.snia.org/ctp/conformingproviders/ibm.html#sftw3
11.2.x	• SMI-S version 1.4 • Test version 1.4.0.24	http://www.snia.org/ctp/conformingproviders/ibm.html#sftw3
11.3.x	• SMI-S version 1.4 • Test version 1.4.0.24	http://www.snia.org/ctp/conformingproviders/ibm.html#sftw3
11.4.x	• SMI-S version 1.4 • Test version 1.4.0.24	http://www.snia.org/ctp/conformingproviders/ibm.html#sftw3
11.5.x	• SMI-S version 1.4 • Test version 1.4.0.24	XIV CIM 11.5.0 has passed CTP1.4 test. Because CTP1.4 test was deprecated from 2013.11.8 by SNIA, the result is not posted on the SNIA website.

XIV CIM Agent release	Official CTP Test Suite	Official CTP Test Results
11.6.x	- SMI-S version 1.6 - Test version 1.6.0.816	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 the IBM XIV Storage System Gen3 with the XIV CIM Agent version 11.6.0.

Publications and related information

Product manuals, other IBM® publications, and websites contain information that relates to the IBM XIV® Storage System.

- [IBM XIV Storage System version 11.6.x publications](#)
The following publications support IBM XIV Storage System software version 11.6.x and its supported hardware models: 281x-114, 281x-214, and 281x-314.
- [IBM XIV Storage System version 11.5.x publications](#)
These product documents support IBM XIV Storage System 11.5.x and XIV system model 281x-114.
- [Warranty, notices, and licensing publications](#)
These publications provide information about warranty, notices, and licensing for IBM XIV Storage System.
- [Related publications and websites](#)
Listed here are the IBM Redbooks® publications, technical papers, and websites that relate to IBM XIV Storage System.

IBM XIV Storage System version 11.6.x publications

The following publications support IBM XIV Storage System software version 11.6.x and its supported hardware models: 281x-114, 281x-214, and 281x-314.

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

For release note publications, see [Release notes PDF publications](#).

IBM XIV Storage System publications

Table 1. IBM XIV Storage System 11.6.x publications

Title	Description	PDF download link
<i>Product Overview</i>	This publication provides introductory and conceptual information about the IBM XIV Storage System Gen3 functional features. Publication number: GC27-3912-10	 Download
<i>Planning Guide</i>	This publication describes pre-installation, networking connectivity, and other requirements and planning considerations for the IBM XIV Storage System model 281x-114, 281x-214, and 281x-314. Publication number: SC27-5412-09	 Download
<i>Command-Line Interface (CLI) Reference Guide</i>	This publication describes the available IBM XIV command-line interface (XCLI) commands, parameter definitions, and return codes for IBM XIV Storage System. Seventh Edition (May 2022)	 Download
<i>Application Programming Interface (API) Reference Guide</i>	This publication describes the IBM XIV Storage System Open Application Programming Interface (API), which is referred to as the Common Information Model (CIM) agent. Publication number: GC27-3916-09	 Download

IBM XIV Management Tools and IBM Hyper-Scale Manager publications

Important: As of June 2022, the IBM XIV Management Tools for XIV Storage System is no longer being maintained. This includes the following : XIV Management Tools 4.8.0.6 and lower, XIV XCLI 4.8.0.6 and lower, and IBM Hyper-Scale Manager v1.9.6 and lower. If these utilities are currently being used, it is important to migrate to the latest IBM Hyper-Scale Manager and XCLI Utility for the most up-to-date security patches. These can both be downloaded from [Fix Central](#). For more information about the XCLI Utility, see [Release notes for IBM XCLI Utility](#) and [Command-line reference](#). For more information about Hyper-Scale Manager, see [Hyper-Scale Manager documentation](#).

Table 2. IBM XIV Management Tools and IBM Hyper-Scale Manager publications

Title	Description	PDF download link
<i>Management Tools Operations Guide</i>	This publication describes how to operate IBM XIV Management Tools 4.8, including the IBM Hyper-Scale Manager and IBM XIV Storage Management GUI tasks. Publication number: SC27-5986-06	 Download

Title	Description	PDF download link
<i>Management Tools XCLI Utility User Guide</i>	This publication describes how to use the XCLI Utility that is provided in the IBM XIV Management Tools 4.x package. Publication number: GC27-3915-05	 Download
<i>IBM Hyper-Scale Manager User Guide</i>	This publication describes how to install and use IBM Hyper-Scale Manager 1.9. Publication number: GC27-5984-05	 Download
<i>IBM Hyper-Scale Manager Quick-Start Guide</i>	This quick guide walks you through installing the IBM Hyper-Scale Manager as an application.	 Download
<i>IBM Hyper-Scale Manager RESTful APIs Specifications</i>	This publication provides specifications of IBM Hyper-Scale Manager Representational State Transfer (REST) APIs version 4.6.	 Download

XIV® Remote Support Proxy publication

Table 3. IBM XIV Remote Support Proxy publication

Title	Description	PDF download link
<i>User Guide</i>	This publication provides installation and user guidance for IBM Remote Support Proxy version 1.3.x. Ninth Edition (July 2022)	 Download

IBM XIV Storage System version 11.5.x publications

These product documents support IBM XIV Storage System 11.5.x and XIV system model 281x-114.

You can also order or download individual publications that have an order number from the IBM [Publications Center](#) website. Publications without an order number can be downloaded here.

For release note publications, see [Release notes PDF publications](#).

Table 1. IBM XIV 11.5.x product publications

Title	Description	Form number
<i>Planning Guide</i>	This publication describes pre-installation and networking connectivity requirements for the IBM XIV Storage System model 281x-114 and 281x-214.	SC27-5412
<i>Product Overview</i>	This publication describes introductory and conceptual material about the IBM XIV Storage System and IBM XIV Storage System software.	GC27-3912
<i>Commands Reference</i>	This publication describes the available IBM XIV command-line interface (XCLI) commands, parameter definitions, and completion code for the IBM XIV Storage System.	GC27-3914
<i>Application Programming Interface Reference</i>	This publication describes the IBM XIV Storage System Open Application Programming Interface (API), which is referred to as the Common Information Model (CIM) agent.	GC27-3916

Table 2. IBM XIV Management Tools 4.5.x product publications

Title	Description	Form number
<i>Management Tools Operations Guide</i>	This publication describes how to operate the XIV management tools, including the IBM Hyper-Scale Manager and IBM XIV Storage Management GUI tasks.	SC27-5986
<i>IBM Hyper-Scale Manager 1.6 Installation as application</i>	This publication describes how to install the IBM Hyper-Scale Manager as an application	GC27-5984
<i>IBM Hyper-Scale Manager 1.6 Installation as a virtual appliance</i>	This publication describes how to install the IBM Hyper-Scale Manager as a virtual appliance	GC27-5985
<i>IBM Hyper Scale Manager 1.6 Installation as Application Quick-Start</i>	This quick guide walks you through installing the IBM Hyper-Scale Manager as an application and using it with the IBM XIV Storage Management GUI.	Quick-start
<i>IBM Hyper Scale Manager 1.6 Installation as a virtual appliance Quick-Start</i>	This quick guide walks you through installing the IBM Hyper-Scale Manager as a Virtual Appliance and using it with the IBM XIV Storage Management GUI.	Quick-start
<i>XCLI Utility User Manual</i>	This publication describes how to use the IBM XIV Storage System command-line interface (XCLI) to run XCLI utility and system commands.	GC27-3915
<i>IBM Hyper-Scale Manager 2.1 RESTful API</i>	This publication provides specifications of IBM Hyper-Scale Manager REST API.	RESTful API

Warranty, notices, and licensing publications

These publications provide information about warranty, notices, and licensing for IBM® XIV® Storage System.

See the *Agreements and License Information* CD that was shipped with the XIV system for the following documents:

- License Information
- Notices and Information
- Supplemental Notices and Information

Table 1. IBM XIV warranty, notices, and licensing publications

Title	Order number
<i>Warranty Information for IBM XIV Storage System</i>	GC27-2268-00

Title	Order number
IBM Safety Notices for IBM XIV Storage System	Environmental and safety notices publications page
IBM Systems Environmental Notices and User Guide	Z125-5823-06
International Agreement for Acquisition of Software Maintenance	Z125-6011-04
License Agreement for Machine Code	SC28-6872-02
International Program License Agreement	Z125-3301-14

Related publications and websites

Listed here are the IBM® Redbooks® publications, technical papers, and websites that relate to IBM XIV® Storage System.

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.

IBM Redbooks publications and technical papers

Various IBM Redbooks publications, Redpapers, and white papers are available for the IBM XIV Storage System. For more papers, see the [IBM XIV Storage System website](#).

Table 1. IBM XIV related publications

Title	Description
IBM XIV Storage System: Architecture, Implementation, and Usage	This Redbooks publication describes the concepts, architecture, and implementation of the XIV system, which is a scalable enterprise storage system that is based on a grid array of hardware components.
IBM XIV Storage System with the Virtual I/O Server and IBM i	This Redbooks publication describes how you can connect the XIV system to the IBM i operating system through the Virtual I/O Server (VIOS). A connection through the VIOS is especially useful for IT centers that have many small IBM i partitions. When you use the VIOS, the Fibre Channel host adapters can be installed in the VIOS and shared by many IBM i clients by using virtual connectivity to the VIOS.
XIV Storage System: Host Attachment and Interoperability	This Redbooks publication describes how to attach an XIV system to various hosting operating systems in combination with databases and other storage-oriented application software. It also provides solutions for combining the IBM XIV Storage System with other storage systems, host servers, or gateways.
IBM XIV Storage System: Copy Services and Migration	This Redbooks publication describes IBM XIV Storage System copy and migration functions for various data protection scenarios to enhance your business continuance, data migration, and online-backup solutions. These include point-in-time copies (also known as snapshots and full volume copies) and remote-copy capabilities in synchronous or asynchronous mode. This book also describes how to integrate the snapshot function with the IBM Tivoli® FlashCopy® Manager, built-in migration capability, and migration alternatives that are based on the IBM SAN Volume Controller.

Related information

View these websites to get more information about the XIV system.

Table 2. IBM XIV related websites

Title	Description	
IBM website	Find more information about IBM products and services.	
IBM Support Portal website	Find support-related information such as downloads, documentation, troubleshooting, and service requests and PMRs.	
IBM Directory of Worldwide Contacts website	Find contact information for general inquiries, technical support, and hardware and software support by country.	
IBM XIV Storage System website	Use this website to learn about the XIV system, including features and hardware summary. This website also has links to white papers, Redbooks publications, and product documentation.	
IBM Systems Storage forum website	Use this forum to share ideas with knowledgeable experts and discover how the latest IBM storage solutions can address your business challenges. Forum topics include storage management, storage virtualization, business continuity, infrastructure simplification, disk storage systems, and storage software products and solutions.	
External information source	Description	PDF download link

External information source	Description	PDF download link
<i>FSTC RFID IT Assets – Phase 2 Business Case and Best Practices</i>	This report contains RFID requirements and specifications, published by the Financial Services Technology Consortium (FSTC).	 Download

Glossary

This glossary includes terms and definitions for IBM® XIV® Storage System.

This glossary includes selected terms and definitions from:

- The *American National Standard Dictionary* for Information Systems, ANSI X3.172–1990, copyright 1990 by the American National Standards Institute (ANSI), 11 West 42nd Street, New York, New York 10036. Definitions derived from this book have the symbol (A) after the definition.
- IBM Terminology, which is available online at the [IBM Terminology website](#). Definitions derived from this source have the symbol (GC) after the definition.
- The *Information Technology Vocabulary* developed by Subcommittee 1, Joint Technical Committee 1, of the International Organization for Standardization and the International Electrotechnical Commission (ISO/IEC JTC1/SC1). Definitions derived from this book have the symbol (I) after the definition. Definitions taken from draft international standards, committee drafts, and working papers that the ISO/IEC JTC1/SC1 is developing have the symbol (T) after the definition, indicating that final agreement has not been reached among the participating National Bodies of SC1.

This glossary uses the following cross-reference forms:

See

Refers the reader to one of two kinds of related information:

- A term that is the expanded form of an abbreviation or acronym. This expanded form of the term contains the full definition.
- A synonym or more preferred term

See also

Refers the reader to one or more related terms.

Contrast with

Refers the reader to a term that has an opposite or substantively different meaning.

A

access

To obtain computing services or data.

In computer security, a specific type of interaction between a subject and an object that results in flow of information from one to the other.

Active Directory

Microsoft Active Directory (AD) provides directory (lookup), DNS and authentication services.

alerting event

An event that triggers recurring event notifications until it is cleared.

allocated storage

The space that is allocated to volumes but not yet assigned. Contrast with *assigned storage*.

API

See *application programming interface (API)*.

application programming interface (API)

An interface that allows an application program that is written in a high-level language to use specific data or functions of the operating system or another program.

assigned storage

The space that is allocated to a volume and that is assigned to a port.

Asynchronous interval

Denotes, per given coupling, how often the master runs a new sync job.

authorization level

The authorization level determines the permitted access level to the various functions of the IBM XIV Storage Management GUI:

Read only

Only viewing is allowed.

Full

Access to all the configuration and control functions is allowed, including shutdown of the system. This level requires a password.

auto-delete priority

As the storage capacity reaches its limits, snapshots are automatically deleted to make more space. The deletion takes place according to the value set for each snapshot, as follows:

1

last to be deleted

4

first to be deleted

Each snapshot is given a default auto delete priority of 1 at creation.

B

basic mode

A means of entering XCLI commands on the XCLI client that requires specifying IP address and login information for each command. Additional output formatting options are available in basic mode.

best effort mode

A mode of remote mirroring in which I/O operation is not suspended when communication between a primary and secondary volume is broken.

C

call home

A communication link established between the storage system and a service provider. The storage product can use this link to call IBM or to another service provider when it requires service. With access to the storage system, service personnel can perform service tasks, such as viewing error logs and problem logs or initiating trace and dump retrievals.

clearing events

The process of stopping the recurring event notification of alerting events.

CLI

The IBM XIV command-line interface (XCLI). See *command-line interface (CLI)*

command-line interface (CLI)

The nongraphical user interface used to interact with the system through set commands and functions. The IBM XIV command-line interface (XCLI) for the IBM XIV Storage System.

completion code

The returned message sent as a result of running CLI commands.

consistency group

A cluster of specific volumes for which a snapshot can be taken simultaneously as a group, thus creating a synchronized snapshot. The volumes in a consistency group are grouped into a single volume set. Snapshots can be taken for the volume set in multiple snapshot sets under the specific consistency group. See also *snapshot set*, *volume set*.

coupling

A primary volume and a secondary volume connected together through mirroring definitions.

D

data availability

The degree to which data is available when needed. Availability is typically measured as a percentage of time in which the system is able to respond to data requests (for example, 99.999% available).

data module

A module dedicated to data storage. A fully populated rack contains nine dedicated data modules, each with 12 disks.

default storage pool

The default storage pool when a volume is created.

destination

See *event destination*.

E

escalation

A process in which event notifications are sent to a wider list of event destinations because the event was not cleared within a certain time.

event destination

An address for sending event notifications.

event notification rule

A rule that determines which users are to be notified, for which events and by what means.

event notification

The process of notifying a user about an event.

event

A user or system activity that is logged (with an appropriate message).

F

fabric

The hardware that connects workstations and servers to storage devices in a SAN. The SAN fabric enables any-server-to-any-storage device connectivity through the use of Fibre Channel switching technology.

FC-AL

Also known as arbitrated loop. A Fibre Channel topology that requires no Fibre Channel switches. Devices are connected in a one-way loop fashion.

FC-HBA

Fibre Channel host bus adapter.

FC

See *Fibre Channel*.

Fibre Channel

Serial data transfer architecture developed by a consortium of computer and mass storage device manufacturers and now being standardized by ANSI.

functional area

One of the high-level groupings of icons (functional modules) of the left pane in the IBM XIV Storage Management GUI screen (for example, Monitor, Configuration, or Volume management). See *functional module*.

functional module

One of the icons of a functional area, on the left pane in the IBM XIV Storage Management GUI screen. For example, System (under Monitor) or Hosts and LUNs (under Configuration). See *functional area*.

G

Graphical user interface (GUI)
On-screen user interface supported by a mouse and a keyboard.

GUI
See *graphical user interface (GUI)*.

H

H/W
Hardware.

HBA
Host bus adapter.

host interface module
The interface data module serves external host requests with the ability to store data. A fully populated rack has six interface data modules.

host
A port name of a host that can connect to the system. The system supports Fibre Channel and iSCSI hosts.

I

I/O
input/output.

image snapshot
A snapshot that has never been unlocked. It is the exact image of the master volume it was copied from, at the time of its creation. See also *snapshot*.

interactive mode
A means of entering XCLI commands on the XCLI client in which the IP address, user, and password information does not need to be specified for each command.

Internet Protocol
Specifies the format of packets (also called datagrams), and their addressing schemes. See also *Transmission Control Protocol (TCP)*.

IOPs
input/output (I/O) per second.

IP
See *Internet Protocol*.

iSCSI
Internet SCSI. An IP-based standard for linking data storage devices over a network and transferring data by carrying SCSI commands over IP networks.

L

latency
Amount of time delay between the moment an operation is issued, and the moment it is committed.

LDAP
Lightweight Directory Access Protocol.

LDAP attribute
A property of an LDAP object, with a single or multiple values. A special object attribute is designated by an LDAP administrator to hold user group memberships values corresponding to XIV roles.

LDAP authentication
A method for authenticating users by validating the submitted credentials against data stored on an LDAP directory.

LDAP directory
A hierarchical database stored on an LDAP server and accessed through LDAP calls.

LDAP server
A server that provides directory services through LDAP.

LDAP status
The status of an LDAP server.

load balancing
Even distribution of load across all components of the system.

locking
Setting a volume (or snapshot) as unwritable (read-only).

LUN map
A table showing the mappings of the volumes to the LUNs.

LUN
Logical unit number. Exports a systems volume into a registered host.

M

main-power cable
The electrical connection between the ac power source and the automatic transfer switch (ATS).

mandatory mode
A mode of remote mirroring in which I/O operation stops whenever there is no communication to the secondary volume.

master volume
A volume that has snapshots is called the master volume of its snapshots.

MIB
Management Information Base. A database of objects that can be monitored by a network management system. SNMP managers use standardized MIB formats to monitor SNMP agents.

Microsoft Active Directory
See *Active Directory*

mirror volume

A volume that contains a backup copy of the original volume.

mirroring
See *remote mirroring*.

modified State
A snapshot state. A snapshot in modified state can never be used for restoring its master volume.

multipathing
Used for direct access from host-interface modules to any volume.

P

peer
Denotes a constituent side of a coupling. Whenever a coupling is defined, a designation is specified for each peer - one peer is designated primary and the other is designated secondary.

pool
See *storage pool*.

primary volume
A volume that is mirrored for backup on a remote storage system.

R

rack
The cabinet that stores all of the hardware components of the system.

remote mirroring
The process of replicating a volume on a remote system.

remote target connectivity
A definition of connectivity between a port set of a remote target and a module on the local storage system.

remote target
An additional storage system used for mirroring, data migration, and so on.

role
The actual role that the peer is fulfilling as a result of a specific condition, either a master or a subordinate.

rule
See *event notification rule*.

S

SAN
Storage area network.

SCSI
Small computer system interface.

secondary volume
A volume that serves as a backup of a primary volume.

Simple Network Monitor Protocol
A protocol for monitoring network devices. See also *MIB*, *SNMP agent*, *SNMP manager*, *SNMP trap*.

SMS gateway
An external server that is used to send SMSs.

SMTP gateway
An external host that is used to relay email messages through the SMTP protocol.

snapshot set
The resulting set of synchronized snapshots of a volume set in a consistency group. See also *consistency group*, *Volume set*.

snapshot
A point-in-time snapshot or copy of a volume. See also *image snapshot*.

SNMP agent
A device that reports information through the SNMP protocol to SNMP managers.

SNMP manager
A host that collects information from SNMP agents through the SNMP protocol.

SNMP trap
An SNMP message sent from the SNMP agent to the SNMP manager, where the sending is initiated by the SNMP agent and not as a response to a message sent from the SNMP manager.

SNMP
See *Simple Network Monitor Protocol*.

snooze
The process of sending recurring event notifications until the events are cleared.

storage pool
A reserved area of virtual disk space serving the storage requirements of the volumes.

Sync Job
A synchronization procedure run by the master at specified user-defined intervals, entailing synchronization between the master and the subordinate.

synchronization
The process of making the primary volume and secondary volume identical after a communication downtime or upon the initialization of the mirroring.

T

target
See *remote target*.

TCP/IP

See *Transmission Control Protocol, Internet Protocol*.

thin provisioning
The ability to define logical volume sizes that are much larger than the physical capacity installed on the system.

Transmission Control Protocol
Transmission Control Protocol (TCP) on top of the Internet Protocol (IP) establishes a virtual connection between a destination and a source over which streams of data can be exchanged. See also *IP*.

trap
See *SNMP trap*.

U

unassociated volume
A volume that is not associated with a consistency group. See *Consistency group*.

uninterruptible power supply
Provides battery backup power for a determined time, so that the system can power down in a controlled manner, on the occurrence of a lengthy power outage.

V

volume cloning
Creating a snapshot from a volume.

volume set
A cluster of specific volumes in a consistency group, for which snapshots are taken simultaneously, thus, creating a synchronized snapshot of all of them. Snapshots of the volume set can be taken into multiple snapshot sets of the specific consistency group. See also *Snapshot set, Volume set*.

volume
A logical address space, having its data content stored on the systems disk drives. A volume can be virtually any size as long as the total allocated storage space of all volumes does not exceed the net capacity of the system. A volume can be exported to an attached host through a LUN. A volume can be exported to multiple hosts simultaneously. See also *Storage pool, Unassociated volume*.

W

WWPN
Worldwide port name

X

XCLI
IBM XIV command-line interface (XCLI) command set. See *command-line interface*.

XCLI client
The system on which the XCLI command is entered.

XCLI identification parameters
Parameters that identify the user issuing the command and the XIV system (if any) on which the command is to run. XCLI identification parameters can be specified:

- By entering them at the beginning of an interactive mode session
- In a configuration file
- When running a command in basic mode
- When running a list of commands as a batch

XCLI system command
An XCLI command that is sent to the XIV system for processing.

XCLI utility command
An XCLI command that is issued on the XCLI client. An XCLI utility command is not sent to an XIV system for processing. XCLI utility commands are used for setting up configurations on the XCLI client and for queries (for example, of software version) that can be processed on the client.

XDRP
The disaster recovery program for the XIV system – The remote mirror feature of the XIV system.

XIV mapping
An association of data on the LDAP server (a specific LDAP attribute) and data on the XIV system. The mapping is used to determine which access rights to grant to an authenticated LDAP user.

IBM XIV 10.2.x publications

These product documents support IBM XIV Storage System 10.2.x and XIV system model 281x-A14.

You can also order or download individual publications that have an order number from the IBM [Publications Center](#) website. Publications without an order number can be downloaded here.

For Release Notes, see [Release notes publications](#).

Table 1. IBM XIV 10.2.x product publications

Title	Description	Order number

Title	Description	Order number
<i>Planning Guide</i>	This publication describes pre-installation and networking connectivity requirements for the IBM XIV Storage System. The <i>Introduction and Planning Guide</i> and the <i>Pre-Installation Network Planning Guide</i> are combined into one publication.	GC27-3913
<i>Product Overview</i>	This publication describes introductory and conceptual material about the IBM XIV Storage System and IBM XIV Storage System software.	GA32-0791
<i>Management Tools User Guide</i>	This publication describes how to install and use XIV management tools, including the IBM Hyper-Scale Manager and IBM XIV Storage Management GUI tasks.	SC27-4230
<i>XCLI Utility User Manual</i>	This publication describes how to use the IBM XIV Storage System command-line interface (XCLI) to start XCLI utility and system commands.	GC27-3915
<i>Commands Reference</i>	This publication describes the available IBM XIV command-line interface (XCLI) commands, parameter definitions, and completion code for the IBM XIV Storage System.	GC27-2213
<i>Application Programming Interface Reference</i>	This publication describes the IBM XIV Storage System Open Application Programming Interface (API), which is referred to as the Common Information Model (CIM) agent.	GC27-3916

- [Release notes publications](#)
This publication contains release notes for the IBM XIV Storage System.

Release notes publications

This publication contains release notes for the IBM® XIV® Storage System.

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

IBM XIV Storage System code

-  [IBM XIV Storage System Version 10.2.4.f-1](#), released on July 2, 2015
-  [IBM XIV Storage System Version 10.2.4.e-8](#), released on April 27, 2015
-  [IBM XIV Storage System Version 10.2.4.e-7](#), released on February 11, 2015
-  [IBM XIV Storage System Version 10.2.4.e-6](#), released on January 14, 2015
-  [IBM XIV Storage System Version 10.2.4.e-5](#), released on October 03, 2013
-  [IBM XIV Storage System Version 10.2.4.e-4](#), released on July 15, 2013
-  [IBM XIV Storage System Version 10.2.4.e-3](#), released on April 10, 2013
-  [IBM XIV Storage System Version 10.2.4.e](#), released on March 6, 2012
-  [IBM XIV Storage System Version 10.2.4.c1](#), released on February 1, 2012
-  [IBM XIV Storage System Version 10.2.4.c](#), released on November 22, 2011

IBM XIV Management Tools

-  [IBM XIV Storage Management GUI Version 4.5.0.2](#)
-  [IBM XIV Storage Management GUI Version 4.5.0.1](#)
-  [IBM XIV Storage Management GUI Version 4.5](#)
-  [IBM XIV Storage Management GUI Version 4.2.2](#)
-  [IBM XIV Storage Management GUI Version 4.2.1](#)
-  [IBM XIV Storage Management GUI Version 4.2](#)
-  [IBM XIV Storage System IBM Hyper-Scale Manager Version 1.3](#)
-  [IBM XIV Storage System IBM Hyper-Scale Manager 1.2](#)
-  [IBM XIV Storage Management GUI version 4.1](#)
-  [IBM XIV Storage System IBM Hyper-Scale Manager 1.1](#)
-  [IBM XIV Storage Management GUI version 4.0](#)
-  [IBM XIV Storage Management GUI version 3.1.1](#)