IBM FlashSystem A9000
Version 12.2

Command-Line Interface (CLI) Reference Guide



SC27-8559-06

Note
Before using this document and the product it supports, read the information in “Notices” on page 711.

Edition notice

Publication number: SC27-8559-06. This publication applies to version 12.2 of IBM FlashSystem A9000 and to all subsequent releases and modifications until otherwise indicated in a newer publication.

© Copyright IBM Corporation 2017.
US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

About this guide ix
Intended audience ix
Conventions used in this guide ix
Related information and publications ix
IBM Publications Center x
Sending or posting your comments x
Getting information, help, and service x

Chapter 1. Overview of the command-line interface (CLI) 1
Overview of the XCLI utility 1
Installing and starting the XCLI 1
Exiting an interactive XCLI session 3
Using the CLI 4
Interactive and basic modes 4
Understanding the command syntax 6
Using identification parameters 7
Display options 11
Running commands in batch mode 12
Displaying CLI help and the XCLI utility version 14

Chapter 2. Host and cluster management commands 15
Adding a host to a cluster 15
Creating a cluster 16
Deleting clusters 17
Listing clusters 18
Removing a host from a cluster 19
Renaming clusters 20
Adding a port to a host 20
Defining a new host 22
Deleting a host 24
Listing hosts 25
Listing ports 26
Removing a port from a host 27
Renaming a host 28
Updating a host definition 29
Mapping a volume to a host or cluster 30
Listing the mapping of volumes to hosts or clusters 33
Setting the special type of hosts or clusters 34
Listing hosts/clusters to which a volume is mapped 35
Unmapping a volume from a host or cluster . . . 36
Setting the default idle time before unmapping a volume 38
Retrieving the default idle time before unmapping a volume 38
Creating a performance class 39
Deleting a performance class 39
Renaming a performance class 41
Listing details of performance classes 41
Adding a host to a performance class 42
Removing a host from its performance class . . . 43
Adding a pool to a performance class 44
Removing a pool from its performance class . . . 45
Adding a volume to a performance class 46
Removing a volume from its performance class . . 47
Adding a domain to a performance class 48
Removing a domain from its performance class . . 49
Setting the rate for a performance class 49
Listing host profiles 51
Updating the host profile 51
Removing the profile of the specified host 52
Enabling the host profiler 53
Disabling the host profiler 53

Chapter 3. Volume management commands 55
Clearing reservations of a volume 55
Listing reservation keys 55
Listing volume reservations 56
Copying volumes 57
Creating a volume 60
Deleting a volume 62
Formatting a volume 64
Listing volumes 65
Listing a volume's extended attributes 67
Locking a volume 68
Renaming a volume 70
Resizing a volume 71
Unlocking a volume 74

Chapter 4. Volume snapshot management commands 77
Changing a snapshot deletion priority 77
Creating a snapshot 78
Deleting a snapshot 81
Duplicating a snapshot 83
Formatting a snapshot 84
Listing snapshot information 86
Restoring a volume from a snapshot 88

Chapter 5. Consistency group management commands 91
Adding a volume to a consistency group 91
Creating consistency groups 93
Deleting a consistency group 94
Listing consistency groups 96
Removing a volume from a consistency group . . . 97
Renaming a consistency group 98
Suspending I/O execution on a consistency group 99
Resuming I/O execution 101
Listing the status of consistency groups with paused I/O 102
Creating a cross-system consistency group 103
Associating an existing consistency group with a cross-system consistency group definition 104
Removing a consistency group from a cross-system consistency group 105

© Copyright IBM Corp. 2017

Adding a remote system name to a cross-system consistency group definition 106
Removing a remote system from a cross-system consistency group 107
Listing cross-system consistency group definitions 108
Retrieving remote systems in a specified cross-system consistency group 109
Deleting a cross-system consistency group. 110
Listing cross-system consistency group definitions 110

Chapter 6. Snapshot set management commands 113
Snapshotting a consistency group. 113
Changing a snapshot group deletion priority . . . 116
Deleting a snapshot group 117
Disbanding a snapshot group 118
Duplicating a snapshot group 119
Formatting a snapshot group 121
Listing snapshot groups 122
Locking a snapshot group 123
Renaming a snapshot group 124
Restoring a consistency group from a snapshot group 125
Unlocking a snapshot group 127
Setting a snapshot group descriptor 128
Returning a snapshot group's descriptor 129

Chapter 7. Storage pool management commands 131
Moving a consistency group between storage pools 131
Changing the pool limitation, performance class, or threshold parameters 133
Changing pool settings for snapshots 135
Creating storage pools 136
Deleting a storage pool 138
Listing storage pools 139
Renaming a storage pool 141
Resizing a storage pool 142
Moving a volume between storage pools 143

Chapter 8. System management commands 147
Displaying the values of configuration parameters 147
Setting configuration parameters 148
Testing the DNS 150
Displaying help 151
Displaying the current maintenance urgency . . . 152
Shutting down the system 152
Listing the operational state 154
Displaying system usage and data reduction statistics 154
Displaying information about effective and physical capacity 155
Displaying information about effective capacity 157
Displaying system capacity thresholds 158
Changing a system capacity threshold 159
Resuming the system's normal operation 160
Displaying the current time 161
Setting the system's time 162
Listing optional time zones. 162
Setting the time zone 163
Printing the current system version 164
Displaying the values of VPD parameters 165
Setting VPD parameters 166
Displaying the system's MIB file 168
Retrieving the electronic license acceptance status 169
Retrieving a fragment of the electronic license file 170
Accepting the electronic license agreement . . . 170
Enabling command auditing 171
Disabling command auditing 172
Displaying the command audit state. 173
Configuring audit servers 173
Checking the command audit state 174
Retrieving the list of Flash VDisks 175
Enabling CIM service. 176
Disabling the CIM service 176
Displaying the CIM service state 177

Chapter 9. Remote target connectivity commands 179
Setting the threshold of a link disruption duration that triggers an event. 179
Updating the target's mirroring configuration . . 180
Activating connectivity to a remote target 181
Deactivating connectivity to a remote target . . . 182
Defining connectivity to a remote target 183
Deleting connectivity to a remote target 185
Listing target connectivity definitions 186
Defining a remote target. 187
Deleting a remote target. 189
Listing remote targets 190
Allowing remote mirroring access 191
Activating a port 192
Adding a new port to a remote target 193
Deactivating a port 194
Deleting a port from a remote system 195
Listing the ports of a remote target 196
Renaming a remote target 197
Updating the target configuration 198
Adding a Quorum Witness to a target 198
Removing a Quorum Witness from a target . . . 200

Chapter 10. Remote mirroring commands 201
Canceling a snapshot mirror (ad hoc sync job) . . 201
Creating a snapshot mirror (ad hoc sync job) . . . 203
Activating mirroring 207
Changing the RPO for local or remote system . . 210
Changing the designation of mirroring peers . . . 212
Changing the mirroring schedule for remote slave peers 214
Changing the role of a mirrored volume 216
Changing a mirroring schedule for local peers . . 219
Creating a mirroring definition 221
Deactivating mirroring 227
Deleting a remote mirroring definition 229
Viewing the mirroring status 231
Obtaining statistics on past sync jobs 234
Switching roles between master and slave. . . . 236
Retrieving RPO thresholds 238

Setting an RPO threshold 239
Changing the interval of a schedule 240
Creating a schedule object 241
Triggering a schedule. 244
Deleting a schedule object 245
Listing a schedule object. 246
Renaming a schedule. 247
Viewing sync job status 247

Chapter 11. HyperSwap commands 249
Creating a HyperSwap relationship 249
Viewing the status of HyperSwap volumes and consistency groups 255
Activating a HyperSwap relationship 258
Deactivating a HyperSwap relationship. 261
Deleting a HyperSwap relationship 262
Switching roles between Master and Slave volumes 265
Changing a peer role in a HyperSwap volume . . 267
Restoring the availability of a Master volume. . . 270
Creating a HyperSwap volume snapshot (ad hoc sync job) 272
Changing the designation of HyperSwap relationship peers 275
Enabling automatic failover in a HyperSwap relationship 277
Disabling automatic failover in a HyperSwap relationship 279
Converting a HyperSwap relationship into a sync mirror 280
Converting a sync mirror into a HyperSwap relationship 281
Creating a new Quorum Witness 284
Listing Quorum Witnesses 286
Updating a Quorum Witness definition. 288
Renaming a Quorum Witness 289
Deleting a Quorum Witness 290
Activating a Quorum Witness 291
Deactivating a Quorum Witness 292
Listing the Quorum Witness connection status . . 293
Getting Quorum Witness information 295

Chapter 12. Data migration commands 297
Activating data migration 297
Deactivating data migration 298
Defining data migration configuration 299
Deleting the data migration process 301
Listing data migration statuses 302
Testing the data migration definition 304

Chapter 13. IBM Hyper-Scale Mobility commands 305
Creating an IBM Hyper-Scale Mobility relation . . 305
Activating a volume migration 309
Deactivating IBM Hyper-Scale Mobility migration 310
Aborting a defined or activated IBM Hyper-Scale Mobility process 311
Moving the IBM Hyper-Scale Mobility source volume to a Proxy state 312
Deleting an IBM Hyper-Scale Mobility relation . . 314
Listing the IBM Hyper-Scale Mobility status . . . 315

Chapter 14. Event handling commands 317
Generating a custom event 317
Generating a CSS product event 317
Defining a new event notification destination. . . 318
Deleting a destination 322
Listing event notification destinations 323
Renaming a destination 324
Testing a destination 325
Updating an event notification destination . . . 327
Adding a destination to a destination group . . . 330
Creating a destination group 331
Updating an event notification destination group 332
Deleting a destination group 333
Listing destination groups 334
Removing a destination from a destination group 335
Renaming a destination group. 336
Clearing alerting events 337
Listing events 338
Listing uncleared alerting events 340
Setting the threshold for event notification . . . 341
Listing thresholds 342
Activating a rule 343
Creating event notification rules 344
Deactivating a rule 347
Deleting event notification rules 348
Listing event notification rules. 349
Renaming event notification rules 350
Updating an event notification rule 351
Defining an SMS gateway 354
Deleting an SMS gateway 356
Listing SMS gateways 357
Prioritizing SMS gateways 357
Renaming an SMS gateway. 359
Updating an SMS gateway 359
Defining a new SMTP gateway 361
Deleting an SMTP gateway. 362
Listing SMTP gateways 363
Prioritizing SMTP gateways 364
Renaming an SMTP gateway 365
Updating an SMTP gateway 366
Generating an XMPNS admin control event . . . 368

Chapter 15. IP configuration commands 369
Creating a new IP interface. 369
Deleting IP interfaces. 370
Listing IP interface configuration 371
Listing IP interface addresses 373
Showing the status and configuration of Ethernet ports 374
Renaming an IP interface 376
Printing the ARP database of an IP interface . . . 377
Testing the traceroute to a remote IP 377
Testing the traceroute to a remote IP 378
Updating an IP interface. 379
Defining a new IPSec connection 382
Updating an existing IPSec connection 383
Removing an existing IPSec connection. 385
Listing IPSec connections 386

Listing IPSec tunnels 386
Connecting to a support center 387
Defining a support center 389
Deleting a support center 390
Disconnecting from a support center 390
Listing support centers 391
Listing the status of all support centers. 392
Defining a support center connection automatically 393
Listing the configuration of the automatic connection to a support center. 394
Creating a new IP access group 395
Removing an address from an IP access group . . 396
Adding a new address to an IP access group . . . 397
Deleting an existing IP access group. 398
Renaming an existing IP access group 398
Listing IP access groups 399
Listing IP access groups 400

Chapter 16. PKI configuration commands 401
Listing PKI items 401
Generating a certificate signing request. 402
Generating a private key and CSR 403
Deleting the PKI content 404
Changing a PKI symbolic name 404
Importing a signed certificate 405
Importing a PKCS#12 certificate 407
Displaying the details of a signed certificate . . . 408
Updating a PKI certificate or services 409

Chapter 17. InfiniBand commands . . 413
Listing the configured InfiniBand ports. 413
Listing data counters for the enabled InfiniBand HCA ports 414
Listing error counters for the enabled InfiniBand HCA ports. 415
Listing the statuses of the enabled InfiniBand HCA ports 416
Listing the configured InfiniBand switches . . . 417
Listing the configured InfiniBand switch management addresses 419
Listing the configured InfiniBand switch firmware versions 420
Listing the configured InfiniBand switch power values 421
Listing the configured InfiniBand switch voltage values 422
Listing the configured InfiniBand switch temperature values 424
Listing the configured InfiniBand switch fan parts 425
Listing the configured InfiniBand switch PSUs . . 426
Listing the configured InfiniBand switch BBUs . . 427
Listing the configured InfiniBand switch fans. . . 429

Chapter 18. Access control commands 431
Adding an access control definition 431
Deleting an access control definition. 432
Listing access control definitions 433
Adding an LDAP server definition 434
Testing an LDAP configuration 436
Listing LDAP configuration parameters 437
Configuring LDAP in the system 438
Listing LDAP servers defined in the system . . . 443
Listing LDAP server users 444
Listing LDAP-based authentication mode 445
Enabling or disabling LDAP-based authentication mode 445
Updating an LDAP server definition 447
Removing an LDAP server definition 448
Launching the ldapsearch utility 449
Defining a new user 451
Deleting a user 453
Adding users to user groups 454
Creating user groups 455
Deleting a user group 457
Listing user groups 458
Removing a user from a user group 459
Renaming user groups 460
Updating a user group 461
Listing users. 462
Renaming users 463
Updating a user definition 465
Creating a new domain 467
Updating a domain definition 469
Renaming a domain 471
Deleting a domain. 472
Listing domains 473
Listing users per domain 475
Listing objects in domains 476
Listing the global domain 477
Attaching an object to a domain 478
Disassociating object from a domain. 479
Associating users to a domain. 481
Removing a user from a domain 483
Adding a pool to a domain. 484
Removing a pool from a domain 485
Moving a pool from one domain to another . . . 487
Setting the domain attribute 488
Setting domain-related policies 489
Displaying domain-related policies 490
Specifying a user associated with IBM Hyper-Scale Manager 491
Retrieving the user associated with the IBM Hyper-Scale Manager. 492
Setting the application administrator's scope of commands 492
Getting the application administrator's scope of commands 493

Chapter 19. Fibre channel and iSCSI configuration and status commands . 495
Discovering FC hosts 495
Changing FC port configuration 496
Listing FC ports 497
Listing FC port tests 499
Starting an FC port test 502
Aborting an FC port test 503
Listing connectivity to hosts 504
Displaying the InfiniBand connectivity status. . . 505

Chapter 20. Flash enclosure maintenance commands 507
Listing Flash enclosure status 507
Listing a Flash canister status 508
Listing a Flash card status 510
Retrieving the list of Flash fans 511
Retrieving the list of Flash IB adapters 512
Retrieving the Flash control connectivity list . . . 513
Retrieving the list of Flash PSUs 514
Retrieving the list of Flash PIBs 515
Retrieving the list of Flash LED cards 516
Listing Flash BBU components 517
Retrieving the serial number of a Flash enclosure 1S 518

Chapter 21. Hardware maintenance commands 521
Monitoring the redistribution process 521
Displaying the system's average power consumption 522
Getting the values for calculating the system's average power consumption 522
Displaying the system's average temperature . . . 523
Enabling XIV Support access 524
Disabling XIV Support access 525
Displaying the XIV Support window 526
Listing system components 527
Listing module configuration 530
Listing the internal temperature of modules . . . 532
Listing boot media devices in the system 534
Viewing vault devices in the system 536
Listing BBUs in the system 538
Listing PSUs in a module 541
Listing compression adapters in the system . . . 542
Listing fans in a module. 543
Listing NICs in the system 545
Listing DIMMs in the modules 547
Listing CPUs in the modules 552
Listing InfiniBand host card adapters in the storage system 554
Listing CNA adapters in the system 555
Listing data disk devices in the system 556
Listing service statuses 558
Listing system components that require service . . 559
Listing trace snapshot on a module 560
Creating a trace snapshot 561

Chapter 22. Statistics commands . . . 563
Getting performance statistics 563
Retrieving usage history. 568

Chapter 23. Metadata commands . . . 571
Setting metadata 571
Deleting metadata 572
Listing metadata 573
Setting user metadata 574
Listing user metadata 575
Deleting user metadata 576

Chapter 24. Encryption enablement and support commands 577
Disabling encryption 577
Enabling encryption 578
Defining a keyserver 581
Removing a keyserver 583
Displaying keyserver status 584
Checking keyserver status 585
Obtaining a new master key 585
Renaming a keyserver 587
Changing keyserver properties 588
Entering a recovery key 589
Generating recovery keys 591
Retrieving the security administrator's recovery key 593
Rekeying the security administrators 594
Displaying recovery key status 596
Recovering key verification. 597
Recovering key share information 599
Finishing the recovery process. 600
Obtaining a new master key 600
Changing the key management scheme. 602
Viewing the key scheme. 603

Chapter 25. Events 605

Chapter 26. Return codes 709

Notices 711
Trademarks 712

About this guide

This guide describes the command-line interface (CLI) commands for IBM® FlashSystem A9000.

Intended audience

This document serves as a reference for system administrators and all IT staff who manage the IBM FlashSystem® A9000 system from the CLI. This document is also a reference for programmers who want to automate storage system commands.

Conventions used in this guide

Command examples and output examples are documented in monospaced font, with a frame around it.

For example:

- **Command:**

```
vol_rename vol=DBVolume new_name=DBVolume1
```

- **Output:**

```
Command completed successfully.
```

Access control refers to the types of user accounts that are allowed to use a specific command.

Return codes are the possible codes that the system can return after a specific command is issued and completed either successfully or with an error.

Related information and publications

You can find additional information and publications related to IBM FlashSystem A9000 on the following information sources:

- IBM FlashSystem A9000 on IBM Knowledge Center (ibm.com/support/knowledgecenter/STJKMM) – on which you can find the following related publications:
  - IBM FlashSystem A9000 – Release Notes
  - IBM FlashSystem A9000 – Product Overview
  - IBM FlashSystem A9000 – Planning Guide
  - IBM FlashSystem A9000 and IBM FlashSystem A9000R – Application Programming Interface (API) Reference Guide
- IBM Flash Storage and Solutions marketing website(ibm.com/systems/storage/flash)
- IBM Storage Redbooks® website (redbooks.ibm.com/portals/storage)
- IBM Hyper-Scale Manager on IBM Knowledge Center (ibm.com/support/knowledgecenter/SSUMNQ)

© Copyright IBM Corp. 2017

IBM Publications Center

The IBM Publications Center is a worldwide central repository for IBM product publications and marketing material.

The IBM Publications Center website (ibm.com/shop/publications/order) offers customized search functions to help you find the publications that you need. You can view or download publications at no charge.

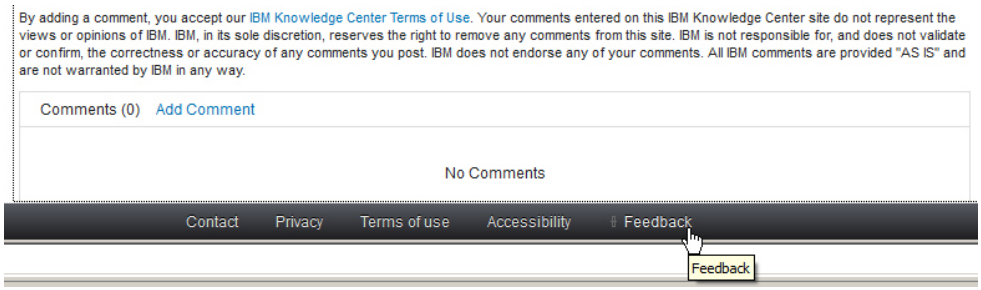
Sending or posting your comments

Your feedback is important in helping to provide the most accurate and highest quality information.

Procedure

To submit any comments about this guide:

- Go to IBM FlashSystem A9000 on IBM Knowledge Center (http://www.ibm.com/support/knowledgecenter/STJKMM), drill down to the relevant page, and then click the **Feedback** link that is located at the bottom of the page.



  The feedback form is displayed and you can use it to enter and submit your comments privately.
- You can post a public comment on the Knowledge Center page that you are viewing, by clicking **Add Comment**. For this option, you must first log in to IBM Knowledge Center with your IBMid.
- You can send your comments by email to starpubs@us.ibm.com. Be sure to include the following information:
  - Exact publication title and product version
  - Publication form number (for example: SC01-0001-01)
  - Page, table, or illustration numbers that you are commenting on
  - A detailed description of any information that should be changed

  Note: When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

Getting information, help, and service

If you need help, service, technical assistance, or want more information about IBM products, you can find various sources to assist you. You can view the following websites to get information about IBM products and services and to find the latest technical information and support.

- IBM website (ibm.com®)

- IBM Support Portal website (www.ibm.com/storage/support)
- IBM Directory of Worldwide Contacts website (www.ibm.com/planetwide)
- IBM developerWorks Answers website (www.developer.ibm.com/answers)
- IBM service requests and PMRs (ibm.com/support/servicerequest/Home.action)

Chapter 1. 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"
- "Using the CLI" on page 4

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"
- "Exiting an interactive XCLI session" on page 3

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 *IBM Hyper-Scale Manager Release Notes* on the IBM Knowledge Center website.

Procedure

Perform these steps to download and install the XCLI:

1. Download the IBM Hyper-Scale Manager installation package from the IBM Fix Central website.
2. Perform one of the following procedures for your operating system.
   - **Windows:** Double-click the installation file, and follow the instructions on the screen.
   - **Linux:** Extract the installation file to a designated folder on your system, using the following command:

     `tar -xzf file_name.tar`
   - **AIX®, HP-UX, Solaris:** Extract the installation file using the following command:

     `gunzip file_name.tar.gz`

     Then extract the file to a designated folder on your system, using the following command:

     `tar -xvf file_name.tar`

© Copyright IBM Corp. 2017

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.

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

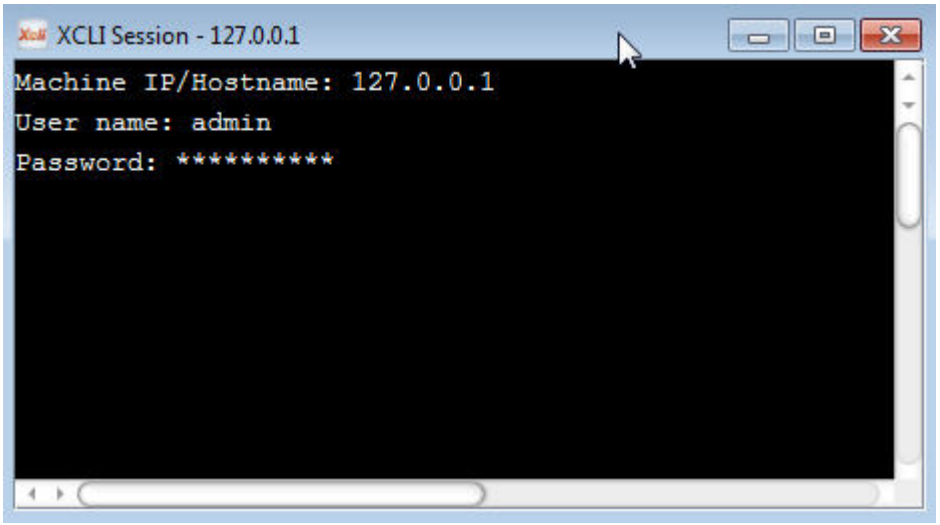


Figure 1. XCLI terminal

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"
- "Understanding the command syntax" on page 6
- "Using identification parameters" on page 7
- "Display options" on page 11
- "Running commands in batch mode" on page 12
- "Displaying CLI help and the XCLI utility version" on page 14

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.

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.

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" on page 8 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" on page 4.

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"
- "Identifying and configuring a storage system" on page 8

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.

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"
- "Table display options" on page 12

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" on page 8 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" on page 4.

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" on page 8 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" on page 4.

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" on page 8 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" on page 151.

The following command displays the XCLI utility version:

```
xcli <-v | --version>
```

Chapter 2. 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.

```
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.
```

© Copyright IBM Corp. 2017

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
[ num_of_visible_targets=num ] >
```

Parameters

Name	Type	Description	Mandatory	
host	Object name	The host name.	Y	

Name	Type	Description	Mandatory	
fcaddress	N/A	FC address of the added port.	N	
iscsi_name	iSCSI initiator name	iSCSI initiator name of the newly added port.	N	
num_of_visible_targets	Integer	Limit the maximum number of target IP addresses to be reported on iSCSI discoveries invoked by this initiator port. Valid values: 0 (unlimited), 2-64.	N	0

The FC port address or iSCSI initiator (port) name assigned to the host must be unique per storage system. The FC port name must be exactly 16 characters long, in hexadecimal format.

Only the following alphanumeric characters are valid: 0-9, A-F, a-f. In addition to the 16 characters, colons (:) may be used as separators in the 16 character port name. The iSCSI initiator name may not exceed 253 characters and may not contain any blank spaces.

Example:

```
host_add_port host=Host1 fcaddress=5001738035C601C0
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- **HOST_BAD_NAME**
  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
- **INVALID_NUM_OF_TARGETS**

  Provided number of visible targets has invalid value. Must be either 0 (unlimited) or number in range 2-64

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

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=tlib_host_pro125_fc0
```

Output:

```
Host                  Type       Port name
--------              ---------  ---------
tlib_host_pro125_fc0  FC         100000062B125CD0
```

Field ID	Field output	Default position
host	Host	1
type	Type	2
port_name	Port Name	3
num_of_visible_targets	iSCSI targets limit	4

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.
- **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

Name	Type	Description	Mandatory	Default
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
- **VOLUME_HAS_INACTIVE_DATA_MIGRATION**
  Cannot map a volume that has an inactive data migration
- **TARGET_NOT_CONNECTED**
  There is currently no connection to the target system
- **REMOTE_TARGET_NOT_CONNECTED**
  There is currently no connection from the target system
- **VOLUME_IS_AN_UNAVAILABLE_HYPERSWAP_PEER**
  Mapping is permitted only on a Hyperswap target that is available for IO.
- **HOST_TYPE_IS_NOT_CONFIGURED**
  Cannot associate a hyperswap volume with a host with unconfigured type. IMPORTANT: Please read the Hyperswap chapter in the 'Best Practice' document to understand the solution requirements.

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
wwn	WWN	6
locked	Locked	7
host	Host	8

Example:

```
mapping_list host=demo__host_1,demo_host_fc10000006072d0190
```

Output:

```
LUN   Volume                                                Proxy   Size   Master
----- ------------------------------------------------------ ------- ------ ------------------
0     vol-2693072-0006                                      no      103
1     vol-2693072-0007                                      no      103
2     cg-2693072-0005.snap_group_00001.vol-2693072-0006     no      103    vol-2693072-0006
3     cg-2693072-0005.snap_group_00001.vol-2693072-0007     no      103    vol-2693072-0007
4     vol-2693172-0013                                      no      103
5     vol-2693172-0013.snapshot_00001                       no      103    vol-2693172-0013

WWN                                Locked   Host
---------------------------------- -------- ---------------------------------------
60017380000035c7000000000000000a   no       tlib_host_host081_fc21000024ff2c4cf7
60017380000035c7000000000000000b   no       tlib_host_host081_fc21000024ff2c4cf7
60017380000035c7000000000000000c   yes      tlib_host_host081_fc21000024ff2c4cf7
60017380000035c7000000000000000c   yes      tlib_host_host081_fc21000024ff2c4cf7
60017380000035c70000000000000011   no       tlib_host_host081_fc21000024ff2c4cf7
60017380000035c70000000000000012   no       tlib_host_host081_fc21000024ff2c4cf7
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Return codes

- **HOST_BAD_NAME**
  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|Windows2008|AllOthers>
```

Parameters

Name	Type	Description	Mandatory
host	Object name	Host name.	N
cluster	Object name	Cluster name.	N
type	Enumeration	Special map type.	Y

Make sure to define the **type** parameter prior to attaching HyperSwap volumes to the host. The supported special types are `hpux`, `zvm`, and `Windows2008`. For any other operating system, select `AllOthers`.

Note: If you need to modify the **type** parameter, make sure to do it when creating a new host definition. Changing the type when volumes are already attached to the host, will cause loss of access to the host.

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

Listing hosts/clusters to which a volume is mapped

Use the **vol_mapping_list** command to list all hosts and clusters to which a volume is mapped.

```
vol_mapping_list vol=VolName
```

Parameters

Name	Type	Description	Mandatory
vol	Object name	Volume name.	Y

This command lists all the hosts and clusters to which a volume is mapped, as well as hosts that are part of a cluster and have host-specific mapping to the volume. The output list contains two columns: name of host/cluster and type (host or cluster).

Field ID	Field output	Default position
host	Host/Cluster	1
type	Type	2
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
force	Boolean	Force completing the unmap operation, even if there exists in-flight I/O that has not completed on the proxy volume.	N	no

The command to unmap from a cluster will unmap the volume from all the hosts that are contained in that cluster.

The command fails if the specified host is contained in a cluster. In this case, the unmapping of the host must be performed through the cluster.

The command does not fail when the volume is not mapped to the host/cluster.

Using this command with **unmap_vol_set_default_idle_time**: The default value of the **idle_seconds** parameter can be set before running the **unmap_volume** command.

The command takes some time to process: If the command fails with **VOLUME_NOT_IDLE** (see the completion codes table below), wait one minute to allow the host to complete background writes, and try 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.

User Category	Permission	Condition
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
- **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

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
type	Enumeration	Determines if associated objects will be limited independently or share the same limit.	N	shared

The performance class name must be unique. Up to 1000 classes can be created.

Example:

```
perf_class_create perf_class=p1
```

Output:

```
Command 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_VOLUMES**

  Deleting Performance Class *Performance Class* will remove the performance limits set for volumes associated with the Performance Class. Are you sure you want to delete Performance Class *Performance Class*?
- **ARE_YOU_SURE_YOU_WANT_TO_DELETE_A_PERF_CLASS_ASSOCIATED_WITH_DOMAINS**

  Deleting Performance Class *Performance Class* will remove the performance limits set for domains associated with the Performance Class. Are you sure you want to delete Performance Class *Performance Class*?

Return codes

- **PERF_CLASS_BAD_NAME**

  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
type	Class type	2
max_iops	Max IO rate (IOPS)	3
max_bw	Max BW 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
- **PERF_CLASS_ASSOCIATED_WITH_POOLS_OR_DOMAINS**
  Performance Class *Performance Class* is already in use by pool or domain.
- **PERF_CLASS_ASSOCIATED_WITH_VOLUMES**
  Performance Class *Performance Class* is already in use by volume.
- **HOST_BAD_NAME**
  Host name does not exist
- **HOST_ALREADY_IN_PERF_CLASS**
  Host *host* already in Performance Class *Performance Class*.

Removing a host from its performance class

Use the **perf_class_remove_host** command to remove a host from its performance class.

```
perf_class_remove_host host=HostName
```

Parameters

Name	Type	Description	Mandatory
host	Object name	The name of the host to be removed from its performance class.	Y

Example:

```
perf_class_remove_host host=h1
```

Output:

```
Command 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
- **PERF_CLASS_ASSOCIATED_WITH_HOSTS**
  Performance Class *Performance Class* is already in use by host.
- **PERF_CLASS_ASSOCIATED_WITH_VOLUMES**
  Performance Class *Performance Class* is already in use by volume.
- **POOL_DOES_NOT_EXIST**
  Storage Pool does not exist
- **POOL_ALREADY_IN_PERF_CLASS**
  Pool *pool name* already in Performance Class *Performance Class*.

Removing a pool from its performance class

Use the **perf_class_remove_pool** command to remove a pool from its performance class.

```
perf_class_remove_pool pool=PoolName
```

Parameters

Name	Type	Description	Mandatory
pool	Object name	The name of the pool to be removed from its performance class.	Y

Example:

```
perf_class_remove_pool pool=h1
```

Output:

```
Command executed 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

- **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 volume to a performance class

Use the **perf_class_add_vol** command to add a volume to a performance class.

```
perf_class_add_vol perf_class=perfClassName vol=VolName
```

Parameters

Name	Type	Description	Mandatory
perf_class	Object name	The name of a performance class.	Y
vol	Object name	The name of the volume to be added to the performance class.	Y

If the volume is already associated with another performance class, it will be removed from that.

Example:

```
perf_class_add_vol perf_class=p1 volume=h1
```

Output:

```
Command 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_ASSOCIATED_WITH_HOSTS**
  Performance Class *Performance Class* is already in use by host.
- **PERF_CLASS_ASSOCIATED_WITH_POOLS_OR_DOMAINS**
  Performance Class *Performance Class* is already in use by pool or domain.
- **VOLUME_BAD_NAME**
  Volume name does not exist
- **VOLUME_ALREADY_IN_PERF_CLASS**
  Volume *volume name* already in Performance Class *Performance Class*.

Removing a volume from its performance class

Use the **perf_class_remove_vol** command to remove a volume from its performance class.

```
perf_class_remove_vol vol=VolName
```

Parameters

Name	Type	Description	Mandatory
vol	Object name	The name of a volume to be removed from its performance class.	Y

Example:

```
perf_class_remove_vol volume=h1
```

Output:

```
Command 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_NOT_CONNECTED_TO_ANY_PERF_CLASS**
  Volume Not connected to any performance class

Adding a domain to a performance class

Use the **perf_class_add_domain** command to add a domain to a performance class.

```
perf_class_add_domain perf_class=perfClassName domain=DomainName
```

Parameters

Name	Type	Description	Mandatory
domain	Object name	The name of the domain to be added to the performance class.	Y
perf_class	Object name	The name of a performance class.	Y

Example:

```
perf_class_add_domain perf_class=perf1 domain=d1
```

Output:

```
Command 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
- **PERF_CLASS_ASSOCIATED_WITH_HOSTS**
  Performance Class *Performance Class* is already in use by host.
- **PERF_CLASS_ASSOCIATED_WITH_VOLUMES**
  Performance Class *Performance Class* is already in use by volume.
- **DOMAIN_ALREADY_IN_PERF_CLASS**
  Domain *domain name* already in Performance Class *Performance Class*.

Removing a domain from its performance class

Use the **perf_class_remove_domain** command to remove a domain from its performance class.

```
perf_class_remove_domain domain=DomainName
```

Parameters

Name	Type	Description	Mandatory
domain	Object name	The name of the domain to be removed from its performance class.	Y

Example:

```
perf_class_remove_domain domain=d1
```

Output:

```
Command 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

Name	Type	Description	Mandatory	Default
max_io_rate	Positive integer	Specifies the performance class maximum rate in IOPS per interface module (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 bandwidth rate per interface module (MB/sec). The maximum setting allowed is 10,000. If zero is specified, the bandwidth rate will not be limited.	N	Keep unchanged.

Either **max_io_rate**, or **max_bw_rate**, or both must be set.

The specified rate is applied to each interface module. To calculate the limit per system, multiply the specified rate by the number of interface modules.

Example:

```
perf_class_set_rate perf_class=p1 max_io_rate=1000
```

Output:

```
Command 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
--------------------------- -------------------- ----------------------
host1                       2012-05-09 22:54:36   Windows 7
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Updating the host profile

Use the **host_profile_set** command to update the host profile.

```
host_profile_set profile_value
```

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.
  Troubleshooting: Try again after the minimal update interval time
- **HOST_BAD_NAME**
  Host name does not exist
- **MAX_HOST_PROFILES_REACHED**
  Maximum number of host profiles already defined

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

Chapter 3. 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.

```
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 ]
```

© Copyright IBM Corp. 2017

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

Copying volumes

Use the **vol_copy** command to copy a source volume onto a target volume.

```
vol_copy vol_src=VolName vol_trg=VolName
```

Parameters

Name	Type	Description	Mandatory
vol_src	Object name	Name of the source volume from which the data is to be taken.	Y
vol_trg	Object name	Name of the target volume to which the data is to be copied.	Y

This command copies a source volume onto a target volume.

All data stored on the target volume is lost and cannot be restored.

This command performs the following as a single atomic action:

- Deletes the target volume.
- Creates a new volume with the same name as the target volume and the same size as the source volume.
- Instantly copies the source volume data onto the target volume.

All volume-to-host mappings of the target volume remain intact during this process. Except for its size, the target volume retains all of its properties, including its name, ID, lock state, creation time and all other attributes.

Immediately after the completion of the command, the volumes are independent of each other and are valid for any further operations (including deletion).

If the target volume is larger then the source volume, excess storage space is freed and returned to the target volume's storage pool. If the target volume is smaller than the source volume, all storage space that is needed to support the additional volume's capacity is reserved from the storage pool.

The command fails in the following cases:

- The target is not formatted.
- The source volume is larger than the target volume, and there is not enough free space in the storage pool that contains the target for target volume resizing.
- The target volume has a snapshot associated with it or if the target volume is a snapshot.
- The target volume is locked.
- The target volume is part of any mirroring definitions (either master or slave).
- The source volume is a slave of a synchronous mirroring, and it is currently inconsistent due to either a re-synchronization or an initialization process.
- There is not enough free space in the storage pool that contains the target.

In the following example, the -y option suppresses the **ARE_YOU_SURE_YOU_WANT_TO_COPY_VOLUME Y/N** prompt.

Example:

```
vol_copy vol_src=DBVolume vol_trg=DBVolumeCopy
```

Output:

```
Command 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*?

Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_COPY_VOLUME**

  Are you sure you want to copy the contents of volume *source Volume* to volume *target Volume*?

Return codes

- **NOT_ENOUGH_SPACE**

  No space to allocate for volume's current usage
- **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_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**

  Volume has snapshots
- **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

- **OPERATION_DENIED_OBJECT_MANAGED**

  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.

- **DATA_REDUCTION_TIER_IS_OFFLINE**

  Data reduced tier is offline, operation is not allowed

  Troubleshooting: Wait for data reduction recovery to start and recover the compression objects

- **SYSTEM_OUT_OF_PHYSICAL_SPACE**

  Operation not allowed while the system is out of physical space.

- **VOLUME_TOO_BIG**

  No space to allocate 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 ]
 [ perf_class=perfClassName ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	Volume name.	Y	N/A
size	Positive integer	Volume size in GB.	N	N/A
size_blocks	Positive integer	Size in number of blocks.	N	N/A
pool	Object name	The name of the storage pool to which the volume belongs.	Y	N/A
ext_id	String	External identifier of the volume.	N	N/A
perf_class	Object name	Name of the performance class for the volume.	N	No performance class

This command is used to create a new volume. The name of the volume must be unique in the system.

The space for the volume is allocated from the specified storage pool and the volume belongs to that storage pool. Specifying the storage pool is mandatory.

When creating a volume, the storage space that is needed to support the volume's capacity is reserved from the capacity of the storage pool for the volume. The command fails if the reservation cannot be committed.

Volumes are created in increments of approximately 1 GB. In some cases, rounding of up to 5% of the size can take place. 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 volume is logically formatted at the creation time, which means that any read operation results in returning all zeros as a response.

Upon successful completion of the command, its lock state is *unlocked*, meaning that write, format and resize operations are allowed.

The creation time of the volume is set to the current time and is never changed.

Example:

```
vol_create vol=DBVolume size=2000 pool=DBPool
```

Output:

```
Command 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
- **VOLUME_TOO_BIG**
  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 offset is illegal
- **OPERATION_DENIED_OBJECT_MANAGED**

This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.

- **ENCRYPTION_IN_PROGRESS**

  System is in the process of changing encryption activation state
- **DOMAIN_MAX_VOLUMES_REACHED**

  The domain exceeds the maximum allowed number of volumes.
- **PERF_CLASS_ASSOCIATED_WITH_HOSTS**

  Performance Class *Performance Class* is already in use by host.
- **PERF_CLASS_ASSOCIATED_WITH_POOLS_OR_DOMAINS**

  Performance Class *Performance Class* is already in use by pool or domain.
- **PERF_CLASS_BAD_NAME**

  Performance Class does not exist
- **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
- **SYSTEM_OUT_OF_PHYSICAL_SPACE**

  Operation not allowed while the system is out of physical space.

Deleting a volume

Use the **vol_delete** command to delete a volume.

```
vol_delete vol=VolName
```

Parameters

Name	Type	Description	Mandatory
vol	Object name	Name of the volume to delete.	Y

After deletion, all data stored on the volume is lost and cannot be restored.

This command cannot be applied to a snapshot. To delete a snapshot, use Deleting a snapshot.

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
- **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_MAPPED_SNAPSHOT**

  Volume which has a snapshot that is mapped to a host cannot be deleted
- **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.
- **VOLUME_HAS_DATA_MIGRATION**

  Data Migration is defined for this volume
- **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

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_DATA_MIGRATION**

  Data Migration is defined for this volume
- **OPERATION_DENIED_OBJECT_MANAGED**

  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  Data reduced tier is offline, operation is not allowed

  Troubleshooting: Wait for data reduction recovery to start and recover the compression objects

Listing volumes

Use the **vol_list** command to list all volumes or a specific one.

```
vol_list [ vol=VolName | pool=PoolName | cg=cgName ] [ show_proxy=<yes|no> ]
[ managed=<yes|no|all> ] [ domain=DomainName ] [ wwn=WWNString ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	Name of a specific volume to be listed.	N	All volumes.
pool	Object name	Name of a specific pool whose volumes are to be listed.	N	Volumes in all Pools.
cg	Object name	List all the volumes in this consistency group.	N	All Consistency Groups.
show_proxy	Boolean	Returns data on proxy volumes (volumes in Proxy state) as well.	N	No
managed	Boolean	Filter only volumes that are or are not managed.	N	no.
domain	Object name	The domain name.	N	All Domains
wwn	String	The WWN in string format.	N	""

This command lists volumes according to:

- Volume name
- Pool
- Consistency Group
- WWN

If no parameter is indicated, the command lists all the available volumes. In addition, the command indicates whether the volume is mirrored.

This command displays the following VAAI fields (available in the XML output format):

- **enable_VAAI**
- **user_disabled_VAAI**

This command displays the following snapshot format field (available in the XML output format):

- **snapshot_format**

Field ID	Field output	Default position
name	Name	1
size	Size (GB)	2
size_MiB	Size (MiB)	N/A
master_name	Master Name	3
cg_name	Consistency Group	4
pool_name	Pool	5
creator	Creator	6
written	Written (GB)	7
written_MiB	Written (MiB)	N/A
proxy	Proxy	N/A
capacity	Capacity (blocks)	N/A
modified	Modified	N/A
sg_name	Snapshot Group Name	N/A
delete_priority	Deletion Priority	N/A
locked	Locked	N/A
snapshot_time	Snapshot Creation Time	N/A
snapshot_time_on_master	Master Copy Creation Time	N/A
snapshot_internal_role	Snapshot Internal Role	N/A
snapshot_of	Snapshot of	N/A
sg_snapshot_of	Snapshot of Snap Group	N/A
wwn	WWN	N/A
mirrored	Mirrored	N/A
locked_by_pool	Locked by Pool	N/A
capacity_used_by_snapshots_MiB	Capacity Used by Snapshots (MiB)	N/A
short_lived_io	Short Live IO	N/A
enable_VAAI	VAAI enabled	N/A
user_disabled_VAAI	VAAI disabled by user	N/A
snapshot_format	Snapshot Format	N/A
unmap_support	Unmap Support	N/A
managed	Managed	N/A

Field ID	Field output	Default position
marked	Marked	N/A
perf_class	Performance Class Name	N/A
thin_provisioning_savings	Thin Provisioning Savings (%)	N/A
est_compression_factor	Est. Compression Factor	N/A
unique_stored_data	Unique Stored Data (GB)	N/A
ha	HA Relation	N/A
target_port_group_id	TPG ID	N/A
target_port_group_state	TPG State	N/A
lock_modes	Lock Modes	N/A

Example:

```
vol_list
```

Output:

```
Name                              Size (GB)   Master Name        Consistency Group
--------------------------------- ----------- ------------------ --------------------
DBLog                             3006
Dev                               2010
Dev.snapshot_00001                2010        Dev

Pool             Creator           Written (GB)
---------------- ----------------- --------------
MainPool         admin             21
MainPool         admin             13
MainPool         admin             0
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Allowed
Read-only users	Allowed
Technicians	Disallowed

Listing a volume's extended attributes

Use the **vol_list_extended** command to return the attributes of the volume which are not returned by **vol_list**.

```
vol_list_extended [ vol=VolName ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	Name of a specific volume to be listed.	N	All volumes.

This command lists extended attributes of volumes according to:

- Volume name

Field ID	Field output	Default position
name	Name	1
wwn	WWN	2
product_serial_number	Product Serial Number	3

Example:

```
vol_list_extended
```

Output:

```
Name                WWN                                Product Serial Number
------------------- ---------------------------------- ------------------------
DBLog               60017380000035C3000000000000000A   MN035C300000000000000A
Dev                 60017380000035C3000000000000000B   MN035C300000000000000B
Dev.snapshot_00001  60017380000035C3000000000000000D   MN035C300000000000000D
Dev.snapshot_00002  60017380000035C3000000000000000E   MN035C300000000000000E
Dev.snapshot_00003  60017380000035C3000000000000000F   MN035C300000000000000F
Marketing           60017380000035C3000000000000000C   MN035C300000000000000C
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Allowed
Read-only users	Allowed
Technicians	Disallowed

Return codes

- **VOLUME_BAD_NAME**

  Volume name does not exist

Locking a volume

Use the **vol_lock** command to lock a volume so that it is read-only.

```
vol_lock vol=VolName
```

Parameters

Name	Type	Description	Mandatory
vol	Object name	Name of the volume to lock.	Y

This command locks a volume so that hosts cannot write to it.

A volume that is locked is write-protected, so that hosts can read the data stored on it, but cannot change it. In addition, a locked volume cannot be formatted or resized. In general, locking a volume prevents any operation (other than deletion) that changes the volume's image.

This command succeeds when the volume's lock state is already set to the one the user is trying to apply. In this case, the lock state remains unchanged.

The lock state of a master volume is set to *unlocked* when a master volume is created.

The lock state of a snapshot is set to *locked* when a snapshot is created.

In addition to the lock state, snapshots also have a modification state. The modification state is a read-only state (which cannot be changed by the user explicitly) and it is initially set to *unmodified* when the snapshot is created. The first time a snapshot lock state is set to *unlocked*, the modification state of the snapshot is changed to *modified*, and it is never changed thereafter.

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.

- **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

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.
- **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

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 rounded up in increments of approximately 1 GB. In some cases, rounding of up to 5% of the size can take place.

If the new size equals the current size, the command succeeds without changes to the volume.

The volume's address space is extended at its end to reflect the increased size, and the additional capacity is logically formatted (that is, zeros are returned for all read commands).

When resizing a regular volume (not a writable snapshot), all storage space that is needed to support the additional volume's capacity is reserved (static allocation). This guarantees the functionality and integrity of the volume, regardless of the resource levels of the volume's storage pool. The command fails if this reservation cannot be committed.

The volume's lock state must be unlocked when the command is issued, or otherwise the command fails.

- Resizing a master volume does not change the size of its associated snapshots.
- These snapshots can still be used to restore their individual master volumes.
- A snapshot is resized in a similar way: the resize does not change the size of its master volume.

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?

Return codes

- **VOLUME_BAD_NAME**
  Volume name does not exist
- **VOLUME_TOO_BIG**
  No space to allocate volume
- **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
- **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_SHRINK_VOLUME**
  Size of volumes cannot be decreased without explicit request
- **VOLUME_SIZE_ABOVE_LIMIT**
  Volume size specified is above limit

- **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

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.

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

User Category	Permission	Condition
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
- **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.
- **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
- **SYSTEM_OUT_OF_PHYSICAL_SPACE**

  Operation not allowed while the system is out of physical space.

Chapter 4. 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.

```
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

© Copyright IBM Corp. 2017

User Category	Permission	Condition
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.
- **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 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

Name	Type	Description	Mandatory	Default
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_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.
- **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.
- **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
- **MAX_SNAPSHOTS_PER_VOLUME_REACHED**
  Maximum number of snapshots already created

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_IS_CONSISTENT_ELCS**

  If a mirrored volume is not consistent then its ELCS is protected and cannot be deleted.
- **OPERATION_DENIED_OBJECT_MANAGED**

This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.

- **DATA_REDUCTION_TIER_IS_OFFLINE**

  Data reduced tier is offline, operation is not allowed

  Troubleshooting: Wait for data reduction recovery to start and recover the compression objects

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
- **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.
- **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
- **MAX_SNAPSHOTS_PER_VOLUME_REACHED**

  Maximum number of snapshots already created

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.

- **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

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
written	Written (GB)	7
written_MiB	Written (MiB)	N/A
proxy	Proxy	N/A
capacity	Capacity (blocks)	N/A
modified	Modified	N/A
sg_name	Snapshot Group Name	N/A
delete_priority	Deletion Priority	N/A
locked	Locked	N/A

Field ID	Field output	Default position
snapshot_time	Snapshot Creation Time	N/A
snapshot_time_on_master	Master Copy Creation Time	N/A
snapshot_internal_role	Snapshot Internal Role	N/A
snapshot_of	Snapshot of	N/A
sg_snapshot_of	Snapshot of Snap Group	N/A
wwn	WWN	N/A
mirrored	Mirrored	N/A
locked_by_pool	Locked by Pool	N/A
capacity_used_by_snapshots_MiB	Capacity Used by Snapshots (MiB)	N/A
short_lived_io	Short Live IO	N/A
enable_VAAI	VAAI enabled	N/A
user_disabled_VAAI	VAAI disabled by user	N/A
snapshot_format	Snapshot Format	N/A
unmap_support	Unmap Support	N/A
managed	Managed	N/A
marked	Marked	N/A
perf_class	Performance Class Name	N/A
thin_provisioning_savings	Thin Provisioning Savings (%)	N/A
est_compression_factor	Est. Compression Factor	N/A
unique_stored_data	Unique Stored Data (GB)	N/A
ha	HA Relation	N/A
target_port_group_id	TPG ID	N/A
target_port_group_state	TPG State	N/A
lock_modes	Lock Modes	N/A

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
- **VOLUME_TOO_BIG**

  No space to allocate volume
- **SNAPSHOT_IS_PART_OF_SNAPSHOT_GROUP**

  Snapshot is part of a Snapshot Group
- **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
- **OPERATION_DENIED_OBJECT_MANAGED**

  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  Data reduced tier is offline, operation is not allowed

Troubleshooting: Wait for data reduction recovery to start and recover the compression objects

Chapter 5. 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.

```
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.

© Copyright IBM Corp. 2017

- 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

Return codes

- **VOLUME_BAD_NAME**

  Volume name does not exist
- **CONS_GROUP_BAD_NAME**

  Consistency Group name does not exist.
- **MAX_VOLUMES_IN_CONS_GROUP_REACHED**

  Consistency Group contains maximum number of volumes.
- **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
- **OPERATION_DENIED_OBJECT_MANAGED**

  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  Data reduced tier is offline, operation is not allowed

  Troubleshooting: Wait for data reduction recovery to start and recover the compression objects

Creating consistency groups

Use the **cg_create** command to create a consistency group.

```
cg_create cg=cgName pool=PoolName
```

Parameters

Name	Type	Description	Mandatory
cg	Object name	Name of the consistency group.	Y
pool	Object name	Storage pool of the consistency group.	Y

This command creates a consistency group. A consistency group is a group of volumes that can all be snapshotted at the same point of time. This is essential for snapshotting several volumes used by the same application or by applications that interact with each other in order to generate a consistent set of snapshots.

The name of the consistency group must be unique in the system. The 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.
- **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

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_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.
- **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

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
mirrored	Mirrored	N/A
ha	HA	N/A
managed	Managed	N/A

Example:

```
cg_list cg=DBgroup
```

Output:

```
Name     Pool Name Mirrored GP Based
DBgroup  default   Yes      No
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Removing a volume from a consistency group

Use the command **cg_remove_vol** to remove a volume from a consistency group.

```
cg_remove_vol vol=VolName
```

Parameters

Name	Type	Description	Mandatory
vol	Object name	Name of the volume to be removed.	Y

This command removes a volume from a consistency group.

A consistency group's name is deduced from the volume name. A unique name is ensured because each volume belongs to only a single consistency group. Future snapshot groups created from this consistency group will not include the snapshot associated with the removed volume.

All the snapshots of the removed volume that were created as part of this consistency group will be permanently removed from the snapshot groups they were associated with.

Following the volume removal:

- The corresponding peer volume is removed from the peer consistency group. If the consistency group is mirrored, the mirroring definition of the removed volume is retained (based on the same settings as the consistency group from which it was removed).
- The peer volume is also removed from the peer consistency group.
- The removed mirrored volume acquires the RPO of the mirrored consistency group from which it was removed.
- An event is generated.

This command succeeds even if the volume is not included in any consistency group.

Requirements for a successful command completion:

- The command can be issued only on the master.
- The link has to be up.
- The consistency group cannot have ongoing sync jobs.

If the command is issued on a mirrored 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
- **OPERATION_DENIED_OBJECT_MANAGED**
  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.
- **DATA_REDUCTION_TIER_IS_OFFLINE**
  Data reduced tier is offline, operation is not allowed
  Troubleshooting: Wait for data reduction recovery to start and recover the compression objects
- **MAX_SNAPSHOTS_PER_VOLUME_REACHED**
  Maximum number of snapshots already created

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

Name	Type	Description	Mandatory
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.
- **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

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.
- **IO_PAUSE_ALREADY_ISSUED_FOR_CONS_GROUP**
  Volume(s) belonging to the Consistency Group are already paused

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.
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

User Category	Permission	Condition
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_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
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 doesnt 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

Chapter 6. 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.

```
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.

© Copyright IBM Corp. 2017

- 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
- **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.
- **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.
- **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.
- **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

- **MAX_SNAPSHOTS_PER_VOLUME_REACHED**

  Maximum number of snapshots already created

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

User Category	Permission	Condition
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.
- **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

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
- **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

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.
- **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

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
- **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.

- **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
- **MAX_SNAPSHOTS_PER_VOLUME_REACHED**

  Maximum number of snapshots already created

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

User Category	Permission	Condition
Technicians	Disallowed	N/A

Return codes

- **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.
- **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

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

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.
- **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

Renaming a snapshot group

Use the **snap_group_rename** command to rename a snapshot group.

```
snap_group_rename snap_group=SnapshotGroupName new_name=Name
```

Parameters

Name	Type	Description	Mandatory
snap_group	Object name	Name of the snapshot group to be renamed.	Y
new_name	Object name	New name for the snapshot group.	Y

Example:

```
snap_group_rename snap_group=DBbackup new_name=DBBackupweekly
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission	Condition
Storage administrator	Allowed	N/A
Storage integration administrator	Allowed	N/A
Application administrator	Conditionally Allowed	At least one of the volumes in the master Consistency Group is mapped to a host or cluster associated with the user and Snapshot Group was created by a server administrator
Security administrator	Disallowed	N/A
Read-only users	Disallowed	N/A
Technicians	Disallowed	N/A

Return codes

- **SNAPSHOT_GROUP_BAD_NAME**

  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.
- **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

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

Name	Type	Description	Mandatory
target_snap_group	Object name	Snapshot group to be restored.	N

Using this command is equivalent to restoring all the volumes in the consistency group, or all the snapshots in the target snapshot group from their snapshots in the snapshot group.

It is possible to restore a snapshot group from a snapshot group.

Requirements for a successful command completion:

- The consistency group or the target snapshot group must contain the exact same volumes that they contained when the snapshot group was generated.
  - Each volume added to the consistency group after the creation of the snapshot group must be removed from the consistency group before restoration is completed.
- The command is inapplicable for a snapshot group of a mirrored consistency group.

See Restoring a volume from a snapshot for more information about the restoring.

Example:

```
snap_group_restore snap_group=DBbackup_copy
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission	Condition
Storage administrator	Allowed	N/A
Storage integration administrator	Allowed	N/A
Application administrator	Conditionally Allowed	Both target and source are snapshots groups of the same master Consistency Group, where at least one of the master volumes in this Consistency Group is mapped to a host or cluster associated with the user, and the target Snapshot Group was created by an application administrator.
Security administrator	Disallowed	N/A
Read-only users	Disallowed	N/A
Technicians	Disallowed	N/A

Return codes

- **VOLUME_HAS_DATA_MIGRATION**

  Data Migration is defined for this volume.
- **SNAPSHOT_GROUP_BAD_NAME**

The snapshot group name does not exist.

- **CONS_GROUP_MISMATCH**
  The snapshot group does not match the consistency group volumes.
- **VOLUME_TOO_BIG**
  No space to allocate to the volume.
- **VOLUME_HAS_MIRROR**
  A mirror is defined for this volume.
- **VOLUME_HAS_HA**
  This operation is forbidden on a volume with a HyperSwap relationship.
- **CONS_GROUP_HAS_MIRROR**
  Mirroring is defined for this consistency group.
- **VOLUME_LOCKED**
  The volume is locked.
- **TARGET_SNAPSHOT_GROUP_BAD_NAME**
  The target snapshot group name does not exist.
- **SNAPSHOT_GROUP_MISMATCH**
  The snapshot group does not match the target snapshot group.
- **TARGET_SNAPSHOT_GROUP_SAME_AS_SOURCE**
  The target snapshot group is identical with the snapshot group.
- **OPERATION_DENIED_OBJECT_MANAGED**
  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.
- **DATA_REDUCTION_TIER_IS_OFFLINE**
  The data reduced tier is offline, the operation is not allowed.
  Troubleshooting: Contact IBM Support

Unlocking a snapshot group

Use the **snap_group_unlock** command to unlock a snapshot group by unlocking all its snapshots.

```
snap_group_unlock snap_group=SnapshotGroupName
```

Parameters

Name	Type	Description	Mandatory
snap_group	Object name	Name of the snapshot group to be unlocked.	Y

This command unlocks a snapshot group by unlocking all its snapshots. This is equivalent to executing Unlocking a volume on each snapshot. Refer to the documentation of Unlocking a volume for a description of unlocking behavior.

Example:

```
snap_group_unlock snap_group=DBbackup
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission	Condition
Storage administrator	Allowed	N/A
Storage integration administrator	Allowed	N/A
Application administrator	Conditionally Allowed	At least one of the volumes in the master Consistency Group is mapped to a host or cluster associated with the user and Snapshot Group was created by a server administrator
Security administrator	Disallowed	N/A
Read-only users	Disallowed	N/A
Technicians	Disallowed	N/A

Return codes

- **SNAPSHOT_GROUP_BAD_NAME**

  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.
- **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

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.
- **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

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.

Chapter 7. 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.

```
cg_move cg=cgName pool=PoolName [ domain_adjust=<yes|no> ]
```

Parameters

Name	Type	Description	Mandatory	Default
cg	Object name	Name of the consistency group to be moved.	Y	N/A
pool	Object name	Name of the target storage pool.	Y	N/A
domain_adjust	Boolean	Adjusts domain resources. If set to True, the resources of the consistency group source domain and destination domain are adjusted to accommodate the consistency group being moved.	N	no

For successful command completion, there must be sufficient space on the target pools.

Example:

```
cg_move cg=DBGroup pool=DBPool
```

Output:

```
Command completed successfully.
```

© Copyright IBM Corp. 2017

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 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_USED_TARGET_NOT_IN_DESTINATION**

  A target that is used by 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.
- **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.

- **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
- **NO_SPACE**

  The system does not have enough free space for the requested Storage Pool size
- **VOLUME_TOO_BIG**

  No space to allocate volume

Changing the pool limitation, performance class, or threshold parameters

Use the **pool_change_config** command to change a storage pool configuration.

```
pool_change_config pool=PoolName [ lock_behavior=<read_only|no_io> ]
[ perf_class=perfClassName ] [ restore_thresholds=<yes|no> | hysteresis=HysteresisValue |
< code=EventCode severity=<INFORMATIONAL|WARNING|MINOR|MAJOR|CRITICAL|NONE>
threshold=<ThresholdValue|NONE> > ]
```

Parameters

Name	Type	Description	Mandatory	Default
pool	Object name	The name of a storage pool.	Y	N/A
lock_behavior	Enumeration	Determines whether and how the pool is locked upon space depletion.	N	read_only
perf_class	Object name	The name of the performance class pool.	N	No performance class
code	N/A	Event code.	N	No code
severity	Enumeration	Severity.	N	No severity
threshold	Integer	The threshold value. None indicates that an event with this severity is not created.	N	No threshold
restore_ thresholds	Boolean	Restore thresholds to default values.	N	no
hysteresis	Integer	The hysteresis of the event throwing.	N	"3"

This command changes the pool behavior when the pool runs out of thin provisioning space.

For thin provisioned storage pools, the **lock_behavior** parameter sets how the pool is locked upon space depletion. The pool can be locked for write, or for both read and write.

Example:

```
pool_change_config pool=VOL_BREED_None_0 lock_behavior=read_only
```

This command changes the Performance Class of the pool.

Example:

```
pool_change_config pool=VOL_BREED_None_1 perf_class=valid_perf_class_name
```

This command changes the thresholds parameters of the pool or reset it to default thresholds value.

Example:

```
pool_change_config pool=VOL_BREED_None_1 code=STORAGE_POOL_VOLUME_USAGE_INCREASED
severity=INFORMATIONAL threshold=40 pool_change_config pool=VOL_BREED_None_1
code=STORAGE_POOL_SNAPSHOT_USAGE_INCREASED severity=INFORMATIONAL threshold=50
pool_change_config pool=VOL_BREED_None_1 restore_thresholds=yes
```

Output:

```
Command 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
- **PERF_CLASS_ASSOCIATED_WITH_HOSTS**
  Performance Class *Performance Class* is already in use by host.
- **PERF_CLASS_ASSOCIATED_WITH_VOLUMES**
  Performance Class *Performance Class* is already in use by volume.
- **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

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 snapshot_size=GB [ lock_behavior=<read_only|no_io> ]
[ perf_class=perfClassName ] [ domain=DomainName ]
```

Parameters

Name	Type	Description	Mandatory	Default
pool	Object name	The name of the new storage pool.	Y	N/A
size	Positive integer	Effective capacity of the storage pool (in gigabytes).	Y	N/A
snapshot_size	Positive integer	Effective capacity allocated for snapshots.	Y	N/A
lock_behavior	Enumeration	Determines whether and how the pool is locked upon space depletion.	N	read_only
perf_class	Object name	The name of the performance class pool.	N	No performance class
domain	Object name	Add the pool to the specified domain.	N	none

The name of the storage pool must be unique in the system. Upon creation, the storage pool is empty and does not contain volumes.

Pool size limits

The parameters **size** and **snapshot_size** relate to effective capacity.

The upper limit of the **size** parameter is set to 1 PB.

As for the lower limits of the **size** and **snapshot_size** parameters, the following restrictions apply:

- **size** cannot be less than 2 TB
- **snapshot_size** must be 0, or not less than 400GB.

These limits can be ignored by force (*-y*). To change the limits, contact the IBM Support.

Example:

```
pool_create pool=DBPool size=1000 snapshot_size=500
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed

User Category	Permission
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- **POOL_SIZE_SMALL**

  Pool size is very small. Volumes may not be able to use this space efficiently. Are you sure?
- **POOL_SNAPSHOT_SIZE_SMALL**

  Pool snapshot size is very small. Snapshots may be deleted frequently. 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
- **PERF_CLASS_ASSOCIATED_WITH_HOSTS**

  Performance Class *Performance Class* is already in use by host.
- **PERF_CLASS_ASSOCIATED_WITH_VOLUMES**

  Performance Class *Performance Class* is already in use by volume.
- **MAX_POOLS_REACHED**

  Maximum number of Storage Pools already defined
- **NO_SPACE**

  The system does not have enough free space for the requested Storage Pool size
- **SNAPSHOT_SIZE_BIGGER_THAN_POOL_SIZE**

  Snapshot size must be equal or smaller than pool size
- **REACHED_POOL_MAX_SIZE**

  Reached max pool size
- **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_CAPACITY_IN_DOMAIN**

  There is not enough free space in the domain.
- **DOMAIN_MAX_POOLS_REACHED**

  The maximum number of domain pools was reached.

Deleting a storage pool

Use the **pool_delete** command to delete a storage pool.

```
pool_delete pool=PoolName
```

Parameters

Name	Type	Description	Mandatory
pool	Object name	The name of the storage pool to be deleted.	Y

This command fails if the storage pool is not empty, that is it still contains volumes.

The capacity of the deleted storage pool is added to the free space.

Example:

```
pool_delete pool=ERPPool
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_DELETE_POOL**

  Are you sure you want to delete Storage Pool *Pool*?

Return codes

- **POOL_DOES_NOT_EXIST**

  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.

Listing storage pools

Use the **pool_list** command to list all storage pools or the specified one.

```
pool_list [ pool=PoolName ] [ managed=<yes|no|all> ] [ domain=DomainName ]
```

Parameters

Name	Type	Description	Mandatory	Default
pool	Object name	The name of a storage pool.	N	All pools.
managed	Boolean	Determines whether to show unmanaged pools (no), managed (yes), or both (all).	N	No
domain	Object name	The domain name.	N	All Domains

When the **pool** parameter is provided, only the specified storage pool is listed.

Example:

```
pool_list
```

Output:

```
Name     Size (GB)  Empty Space (GB)
default  24292      9225
DBPool   1013       1013
```

Field ID	Field output	Default position
name	Name	1
size	Size (GB)	2
size_MiB	Size (MiB)	N/A
snapshot_size	Snap Size (GB)	3
snapshot_size_MiB	Snap Size (MiB)	N/A
total_volume_size	Total Vols (GB)	4
total_volume_size_MiB	Total Vols (MiB)	N/A
empty_space	Empty (GB)	5
empty_space_MiB	Empty (MiB)	N/A
used_by_volumes	Used by Vols (GB)	6
used_by_volumes_MiB	Used by Vols (MiB)	N/A
used_by_snapshots	Used by Snaps (GB)	7
used_by_snapshots_MiB	Used by Snaps (MiB)	N/A
creator	Creator	N/A
locked	Locked	8
lock_behavior	Lock Behavior	N/A
create_last_consistent_ snapshot	Create Last Consistent Snapshot	N/A
protected_snapshot_priority	Protected Snapshots Priority	N/A
managed	Managed	N/A
perf_class	Perf Class Name	9
domain	Domain	10
sparse	Sparse	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Renaming a storage pool

Use the **pool_rename** command to rename the specified storage pool.

```
pool_rename pool=PoolName new_name=Name
```

Parameters

Name	Type	Description	Mandatory
pool	Object name	The current name of the storage pool.	Y
new_name	Object name	The new name of the storage pool.	Y

The new name of the storage pool must be unique in the system.

This command succeeds even if the new name is identical with the current name.

Example:

```
pool_rename pool=DBPool new_name=ERPPool
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- **POOL_DOES_NOT_EXIST**
  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
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.

- When increasing a storage pool size, the command succeeds only if the free space holds enough free capacity to allow such an increase.
- When decreasing a storage pool size, the command succeeds only if the storage pool itself holds enough free capacity to allow such a reduction.
- If the new size equals the current size, the command succeeds without changing the storage pool.

Pool size limits

The parameters **size** and **snapshot_size** relate to effective capacity.

The upper limit of the **size** parameter is set to 1 PB.

As for the lower limits of the **size** and **snapshot_size** parameters, the following restrictions apply:

- **size** cannot be less than 2 TB
- **snapshot_size** must be 0, or not less than 400GB.

These limits can be ignored by force (*-y*). To change the limits, contact the IBM Support.

This command fails if the current storage pool size cannot be decreased, or if the free space 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_SIZE_SMALL**

  Pool size is very small. Volumes may not be able to use this space efficiently. Are you sure?
- **POOL_SNAPSHOT_SIZE_SMALL**

  Pool snapshot size is very small. Snapshots may be deleted frequently. Are you sure?

Return codes

- **POOL_DOES_NOT_EXIST**

  Storage Pool does not exist
- **POOL_SIZE_TOO_SMALL**

  Storage Pool usage exceeds requested size
- **REACHED_POOL_MAX_SIZE**

  Reached max pool size
- **NO_SPACE**

  The system does not have enough free space for the requested Storage Pool size
- **POOL_SNAPSHOT_SIZE_TOO_SMALL**

  Storage Pool snapshot usage exceeds requested snapshot size
- **SNAPSHOT_SIZE_BIGGER_THAN_POOL_SIZE**

  Snapshot size must be equal or smaller than pool size
- **OPERATION_DENIED_OBJECT_MANAGED**

  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.
- **NO_FREE_CAPACITY_IN_DOMAIN**

  There is not enough free space in the domain.

Moving a volume between storage pools

Use the **vol_move** command to move a volume and all its snapshot from one storage pool to another.

```
vol_move vol=VolName pool=PoolName [ domain_adjust=<yes|no> ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	Name of the volume to move.	Y	N/A
pool	Object name	Name of the storage pool to which to move.	Y	N/A
domain_adjust	Boolean	Adjust domain resources. If set to *true*, the resources of the volume source domain and destination domain are adjusted to accommodate the volume being moved.	N	no

When moving a master volume from one storage pool to another, all of its snapshots are moved together with it to the destination storage pool.

This command fails when trying to move a snapshot of a volume on its own. This command can fail due to the lack of either soft or hard space.

The command succeeds only if the destination storage pool has enough free storage capacity to accommodate the volume and its snapshots. The exact amount of storage capacity allocated from the destination storage pool is released at the source storage pool.

A volume which belongs to a consistency group cannot be moved without the entire consistency group. You may use Moving a consistency group between storage pools to move the consistency group itself from one storage pool to another.

A volume that is asynchronously mirrored cannot be moved into a thin provisioning pool.

Example:

```
vol_move vol=DBLog pool=DBPool
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- **VOLUME_BAD_NAME**

  Volume name does not exist
- **POOL_DOES_NOT_EXIST**

  Storage Pool does not exist
- **NOT_ENOUGH_SPACE**

  No space to allocate for volume's current usage
- **NO_SPACE**

  The system does not have enough free space for the requested Storage Pool size
- **VOLUME_TOO_BIG**

  No space to allocate volume
- **NO_FREE_CAPACITY_IN_DOMAIN**

  There is not enough free space in the domain.
- **VOLUME_IS_SNAPSHOT**

  Operation is not permitted on snapshots
- **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_TARGET_NOT_IN_DESTINATION**

  A target that is used by mirror in the pool is not associated with the target domain.
- **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_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
- **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

Chapter 8. System management commands

This section describes the command-line interface (CLI) for system management.

Displaying the values of configuration parameters

Use the **config_get** command to show the values of configuration parameters.

```
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**

© Copyright IBM Corp. 2017

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
Read-only users	Disallowed
Technicians	Allowed

Return codes

- **UNRECOGNIZED_CONFIG_PARAMETER**

  Unrecognized configuration parameter: '*name*'.

  Troubleshooting: Use a valid configuration parameter as an input.
- **READ_ONLY_CONFIG_PARAMETER**

  Configuration parameter: '*name*' is read-only.

  Troubleshooting: You cannot modify read-only parameters.
- **IPV4_NOT_CONFIGURED**

  IPv4 address is not configured on management interface

  Troubleshooting: Define IPv4 addrress for management before disabling IPv6
- **RULE_WITH_SNMP_DEST_EXISTS**

  Can not set snmp_type to none. There is a rule that contains SNMP destination.

Testing the DNS

Use the **dns_test** command to test the DNS (Domain Naming Service).

```
dns_test name=Name [ type=<A|AAAA> ]
```

Parameters

Name	Description	Mandatory	Default
name	Name of the host to be resolved.	Y	N/A
type	Type of query.	N	According to the DNS server type

This command attempts to translate the DNS name into an IP address. Translation is attempted through each of the defined DNS servers.

This command fails if no DNS servers are defined. A failure of the translation from a name to an IP address is not considered a failure of the command.

The result of each defined DNS server is displayed.

Field ID	Field output	Default position
name	Name	1
primary_ip	IP (Primary DNS)	2
secondary_ip	IP (Secondary DNS)	3

Example:

```
dns_test name=hermes.xiv
```

Output:

```
Name         IP (Primary DNS)   IP (Secondary DNS)
------------ ------------------ --------------------
hermes.xiv   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

Field ID	Field output	Default position
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> ] [ ignore_ha=<yes|no> ]
```

Parameters

Name	Type	Description	Mandatory	Default
emergency	Boolean	Instructs the system to shut down within a timeout even if some of the disks could not be saved, much like in an emergency shutdown performed when the system loses power.	N	no
ignore_ha	Boolean	Ignore activated HA objects.	N	no

The system stops serving hosts, de-stages all information to disks and then turns itself off. If the **emergency** parameter is defined, the system shuts down within the timeout period.

NOTE: USING THIS OPTION MAY CAUSE DATA LOSS.

Example:

```
shutdown -y
```

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 IBM Support

- **FLASH_CCL_IN_PROGRESS**

  The requested command cannot be invoked while Flash Enclosure CCL is in progress.

  Troubleshooting: Wait for Flash CCL to complete
- **SYSTEM_HAS_ACTIVE_MASTER_HA_RELATIONS**

  There are ha relations with Master role on the system. If you continue with shutdown without first handling those relations, the host may lose access to those volumes.

  Troubleshooting: The recommended action is to use switch_roles before continuing. You may explicitly add ignore_ha=yes to force the operation.

Listing the operational state

Use the **state_list** command to display the current operational state of the system.

```
state_list
```

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 system usage and data reduction statistics

Use the **system_usage_get** command to retrieve system-wide usage and data reduction related statistics.

```
system_usage_get
```

The command provides various information on system usage, and on data reduction, including:

- the sum of all user volume and snapshot sizes, excluding internal volumes (statistics and metadata)
- the sum of all logical block address (LBA) ranges written to the currently existing volumes and snapshots
- savings due to thin provisioning
- savings due to data reduction
- deduplication and compression factors.

The data retrieved by the command may vary due to currently running background processes.

Example:

```
system_usage_get
```

Output:

```
Volumes and Snapshots (GB)   Thin Provisioning Savings (%)   Total Written (GB)
---------------------------- ------------------------------ --------------------
3100                         90                             340


Data Reduction Savings (%)   Total Stored (GB)   Deduplication Factor   Compression Factor
---------------------------- ------------------ ---------------------- ------------------
87                           45                 1.37                   5.63
```

Field ID	Field output	Default position
total_volumes_and_snapshots	Volumes and Snapshots (GB)	1
thin_provisioning_savings	Thin Provisioning Savings (%)	2
total_written	Total Written (GB)	3
total_written_pending_deletion	Total Written Pending Deletion (GB)	4
data_reduction_savings	Data Reduction Savings (%)	5
total_stored	Total Stored (GB)	6
deduplication_factor	Deduplication Factor	7
compression_factor	Compression Factor	8
data_only_deduplication_factor	Data Only Deduplication Factor	N/A
data_only_compression_factor	Data Only Compression Factor	N/A
data_only_reduction_factor	Data Only Reduction Factor	N/A
deduplication_factor_full_accuracy	Deduplication Factor Full Accuracy	N/A
compression_factor_full_accuracy	Compression Factor Full Accuracy	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Displaying information about effective and physical capacity

Use the **system_capacity_list** command to display information about the system's effective and physical capacity.

```
system_capacity_list
```

The command output displays two lines of information: one for the effective capacity, and one for the physical capacity.

Effective capacity

Allocated effective capacity represents the sum of all virtual capacity provisioned to domains and to the pools in the default domain before any data reduction is applied.

There is a limit to the amount of effective capacity supported by each system. This limit is determined by the system configuration, namely, by the amount of grid controller resources, and it significantly exceeds the system physical capacity.

The command output for effective capacity includes the following field:

- **Allocated (GB)** — The sum of all effective capacity provisioned to domains, and to pools in the default domain. The value is represented in GB.

The rest of the output fields for effective capacity - **Total**, **Allocated (%)**, **Free (GB)**, and **Free (%)** - are currently not available (N/A).

Physical capacity

Physical capacity represents the amount of data that can be stored by the system after data reduction is applied. It is derived from the amount of flash storage media available in the system after taking into account the RAID protection scheme.

The command output for physical capacity includes the following fields:

- **Total** — The system's total physical capacity in GB.
- **Allocated (GB)** — The consumed physical capacity, represented in GB.
- **Allocated (%)** — The allocated capacity, represented as percentage of the total capacity.
- **Free (GB)** — The difference between the total and allocated capacity in GB.
- **Free (%)** — The free capacity, represented as percentage of the total capacity.

Example:

```
system_capacity_list
```

Output:

```
Type        Total (GB)   Allocated (GB) Allocated (%)  Free (GB)  Free (%)
----------- ----------- --------------- -------------- ---------- ---------
Effective   N/A         50640           N/A            N/A        N/A
Physical    16530       1400            8              15130      92
```

Field ID	Field output	Default position
type	Type	1
total	Total (GB)	2
allocated	Allocated (GB)	3
allocated_percent	Allocated (%)	4
free	Free (GB)	5
free_percent	Free (%)	6

Field ID	Field output	Default position
total_MiB	Total (MiB)	N/A
allocated_MiB	Allocated (MiB)	N/A
free_MiB	Free (MiB)	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Displaying information about effective capacity

Use the **system_effective_capacity_get** command to display additional information about the system's effective capacity.

```
system_effective_capacity_get
```

This command displays information about the system's effective capacity limit. This value depends on the current system configuration.

Example:

```
system_effective_capacity_get
```

Output:

```
Limit (GB)
------------
1400071
```

Field ID	Field output	Default position
effective_capacity_limit	Limit (GB)	1
effective_capacity	Effective (GB)	N/A
effective_capacity_factor	Effective Factor	N/A
effective_capacity_max_limit	Max. Limit (GB)	N/A
effective_capacity_default_limit	Default Limit (GB)	N/A
effective_capacity_min_limit	Min. Limit (GB)	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed

User Category	Permission
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Displaying system capacity thresholds

Use the **system_capacity_threshold_list** command to list the current system capacity thresholds.

```
system_capacity_threshold_list
```

A user can define up to 8 progressive thresholds that will trigger events about physical capacity consumption, as per the **system_capacity_list** command (see "Displaying information about effective and physical capacity" on page 155). An event of the configured severity is emitted once if the configured threshold value is exceeded, and an informational event is emitted when capacity recedes below the threshold minus the hysteresis.

For example, if a threshold is set at 85% and the hysteresis is set at 3%:

- a **SYSTEM_CAPACITY_USAGE_INCREASED** event is emitted when system capacity used in percent (per **system_capacity_list**) moves from a value below 85% to a value of 85% or more
- the matching **SYSTEM_CAPACITY_USAGE_DECREASED** event is emitted when system capacity used moves from a value greater than 82% down to a value of 82% or less.

The output includes the following fields:

- The ordinal of the threshold (between 1 and 8)
- Threshold value in percent (between 10 and 99)
- Severity of the event
- Hysteresis value (between 1 and 10, same for all thresholds)
- Indication whether the threshold is enabled or not

Example:

```
system_capacity_threshold_list
```

Output:

```
Ordinal   Threshold (%)   Severity   Hysteresis (%)   Enabled
--------- --------------- ---------- ---------------- ---------
1         70              Warning    3                yes
2         75              Minor      3                yes
3         80              Minor      3                yes
4         85              Major      3                yes
5         90              Major      3                yes
6         95              Critical   3                yes
7         97              Critical   3                yes
8         99              Critical   3                yes
```

Field ID	Field output	Default position
ordinal	Ordinal	1
threshold	Threshold (%)	2
severity	Severity	3
hysteresis	Hysteresis (%)	4
enabled	Enabled	5

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Allowed

Changing a system capacity threshold

Use the **system_capacity_threshold_change** command to change a system capacity threshold.

```
system_capacity_threshold_change hysteresis=HysteresisValue |  < ordinal=Ordinal [ enabled=<
yes|no> ] [ threshold=ThresholdValue ] [ severity=<WARNING|MINOR|MAJOR|CRITICAL> ] >
```

The user can define progressive thresholds that will trigger events about physical capacity consumption. An event of the configured severity is issued once if the configured threshold value is exceeded, and an informational all-clear event is issued, when capacity drops below the threshold minus the hysteresis.

Parameters:

Name	Type	Description	Mandatory	Default
hysteresis	Integer	The hysteresis value in per cent (same for all thresholds).	N	N/A
ordinal	Integer	The ordinal of the threshold.	N	N/A
enabled	Boolean	Enable or disable system capacity threshold.	N	yes
threshold	Integer	The new threshold value in percent, strictly monotonically increasing across thresholds.	N	No threshold
severity	N/A	The new severity value, strictly monotonically increasing across thresholds.	N	No severity

Example:

```
system_capacity_threshold_change hysteresis=3
```

This command changes the hysteresis of the system capacity.

Example:

```
system_capacity_threshold_change ordinal=1 enabled=yes threshold=15 severity=WARNING
```

This command enables ordinal 1 and changes the threshold and severity parameters of the system capacity.

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- **SYSTEM_CAPACITY_DUPLICATE_THRESHOLD**
  New threshold value is duplicating another threshold.
- **SYSTEM_CAPACITY_THRESHOLD_NOT_MONOTONIC**
  Threshold value or severity must be monotonic.

Resuming the system's normal operation

Use the **system_resume_normal_operation** command to move a system back to read-write state after it ran out of physical space.

```
system_resume_normal_operation
```

Example:

```
system_resume_normal_operation
```

Output:

```
Command executed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed

User Category	Permission
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_RESUME_NORMAL_OPERATION**

  The system is still very close to full. Are you sure you want to resume normal operation?

Return codes

- **SYSTEM_IS_STILL_OUT_OF_PHYSICAL_SPACE**

  The system is still out of physical space, normal operation cannot be resumed.
- **SYSTEM_ALREADY_OPERATING_NORMALLY**

  The system is already operating normally.

Displaying the current time

Use the **time_list** command to display the current system time.

```
time_list
```

This command shows the current time, date and time zone.

Field ID	Field output	Default position
time	Time	1
date	Date	2
timezone	Time Zone	3
dst	Daylight Saving Time	4

Example:

```
time_list
```

Output:

```
Time       Date         Time Zone        Daylight Saving Time
---------- ------------ ---------------- ----------------------
10:09:47   2008-02-19   Asia/Jerusalem   no
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Setting the system's time

Use the **time_set** command to set the system's time in YYYY-MM-DD.HH:MM:SS format.

```
time_set time=Timestamp
```

Parameters

Name	Description	Mandatory
time	New current time.	Y

Example:

```
time_set time=2016-03-04.03:02:01
```

Output:

```
Command 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.
- **FLASH_ENCLOSURE_TIME_UPDATE_FAILED**
  Failed to update flash enclosure time
  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	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 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

In configurations that use IBM Netcool® Network Management for managing equipment, an enterprise (private) SNMP MIB from UC Davis is required. This MIB file can be downloaded from: http://www.net-snmp.org/docs/mibs/UCD-SNMP-MIB.txt.

After obtaining the a9000.mib file from the device with the **mib_get** command, note the following **IMPORTS** declaration line, which requires the parent **UCD-SNMP-MIB**:

```
IMPORTS
    MODULE-IDENTITY, OBJECT-TYPE,
        NOTIFICATION-TYPE,
        Gauge32, Integer32  FROM SNMPv2-SMI
                ucdavis FROM UCD-SNMP-MIB
    OBJECT-GROUP, NOTIFICATION-GROUP,
        MODULE-COMPLIANCE    FROM SNMPv2-CONF
        sysName           FROM SNMPv2-MIB

    TEXTUAL-CONVENTION, DisplayString
                              FROM SNMPv2-TC;
```

When both MIB files (a9000.mib and UCD-SNMP-MIB) are imported into MIB Manager, the full OID path is properly defined by the combination of the declarations in each.

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
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="425a68393141592653 59ba94ca1106dd587f84fe"/>
<fragment_size value="20"/>
```

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

- **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. For the instructions on retrieving the correct electronic license version, see below.	Y	N/A
approver_name	String	The approver's name.	N	none

To retrieve the correct electronic license version, proceed as follows:

1. Run the command **elicense_status_get -x**.
2. In the command output, find the string **version value** and copy its value.

Example:

```
elicense_accept version=xiv_license_v11.6.2_with_ela approver_name=johndoe
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- **ELICENSE_INCOMPATIBLE_VERSION**

  The accepted version of the Electronic license dose not match the current version

  Troubleshooting: Please retrieve the current electronic license version and accept it
- **ELICENSE_ALREADY_ACCEPTED**

  Electronic license already accepted

  Troubleshooting: You do not need to accept the electronic license
- **ELICENSE_DISABLED**

  Electronic license check is disabled

  Troubleshooting: You do not need to accept the electronic license

Enabling command auditing

Use the **audit_enable** command to enable CLI command auditing

```
audit_enable
```

This command is used by a security administrator to enable the auditing of user-entered CLI commands on an external auditing server. For this command to

complete successfully, the current auditing state must be DISABLED (that is, the **audit_show** command returns a *no*), and at least one audit server must be configured successfully by the **audit_config_set** command.

Example:

```
xcli -u -c XIV1 audit_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

- **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	Allowed
Storage integration administrator	Disallowed

User Category	Permission
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	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Checking the command audit state

Use the **audit_config_get** command to show the current configuration of CLI command auditing.

```
audit_config_get
```

Field ID	Field output	Default position
primary_server	Primary Server	1
primary_port	Primary Port	2
secondary_server	Secondary Server	3
secondary_port	Secondary Port	4
audit_protocol	Protocol	5

Example:

```
audit_config_get
```

Output:

```
Primary Server   Primary Port   Secondary Server   Secondary Port   Protocol
--------------- -------------- ------------------ ---------------- ----------
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

Retrieving the list of Flash VDisks

Use the **flash_vdisk_list** command to retrieve the list of Flash VDisks.

```
flash_vdisk_list
```

Example:

```
flash_vdisk_list
```

Output:

```
Name            Enclosure Id          Status   ID   Lun
--------------- --------------------- -------- ---- ------------------
xiv_vdisk_2_0   1:Flash_Enclosure:2   OK       0    0000000000000000

Capacity(bytes)   Block Size
----------------- ------------
5717176090624     512
```

Field ID	Field output	Default position
name	Name	1
enclosure_id	Enclosure Id	2
status	Raid Status	3
vdisk_id	ID	4
lun	Lun	5
capacity_in_gb	Capacity(GB)	6
block_size	Block Size(bytes)	7
capacity_in_gib	Capacity(GiB)	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed

User Category	Permission
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Enabling CIM service

Use the **cim_enable** command to enable the CIM service.

```
cim_enable
```

This command enables the CIM service and the associated SLP service. In order for this command to complete successfully, the current CIM service state must be DISABLED (that is, the **cim_show** command returns *no*).

Example:

```
cim_enable
```

Output:

```
Command 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

Chapter 9. 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.

```
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**

© Copyright IBM Corp. 2017

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
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

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
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

Listing target connectivity definitions

Use the **target_connectivity_list** command to list all the connectivity definitions of a remote target.

```
target_connectivity_list [ target=TargetName ]  [ domain=DomainName ]
```

Parameters

Name	Type	Description	Mandatory	Default
target	Object name	Target name that is listed.	N	All targets
domain	Object name	The domain name.	N	All Domains

Field ID	Field output	Default position
target_name	Target Name	1
remote_port_address	Remote Port	2
local_fc_port	FC Port	3
local_ip_port	IP Interface	4
active	Active	5
up	Up	6

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Defining a remote target

Use the **target_define** command to define a new remote target for remote mirroring or data migration.

```
target_define target=TargetName  protocol=<FC|iSCSI> [ iscsi_name=iSCSIName ]
 [ xiv_features=<yes|no> ]  [ system_id=SystemId ] [ domain=DomainList ]
 [ quorum_witness=QW_Name ] [ uses_512b_sectors=<yes|no> ]
```

Parameters

Name	Type	Description	Mandatory	Default
target	Object name	Local name of the remote target.	Y	N/A
protocol	Enumeration	FC (Fiber Channel) or iSCSI, depending on the communication protocol supported by the remote host.	Y	N/A
iscsi_name	iSCSI initiator name	iSCSI name of the remote target. This field is mandatory for iSCSI hosts.	N	N/A
system_id	String	ID of the remote system. Should be the same as the output of 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
quorum_witness	Object name	The name of the quorum witness that is associated with the target.	N	none
uses_512b_sectors	Boolean	Optimize the asynchronous mirror data transfer for remote targets with 512B sector size.	N	No

This command defines the communication topology between a local storage system and a remote storage system to enable various features, such as remote mirroring. The local storage system can write to or read from the remote storage system, or allow the target storage system to write to or read from it.

The first step when defining a new target connectivity is to specify the name of the remote storage system and the protocol used to communicate with it. There are two possible protocols: Fiber Channel (FC) and iSCSI. Each remote target is available through only one of these protocols.

This step only defines the remote system object. No connectivity definitions are defined yet and no communications are performed yet.

Once you have defined a remote target, the only way to change its protocol type is to delete the remote target and define it again.

Example:

```
target_define target=Nextra2 protocol=FC
```

Output:

```
Command 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.
- **QUORUM_WITNESS_BAD_NAME**
  Quorum witness name does not exist.
- **QUORUM_WITNESS_IS_NOT_ACTIVATED**
  The quorum witness is not activated.
- **QUORUM_WITNESS_CANNOT_BE_ADDED_TO_A_TARGET_OF_THIS_TYPE**
  A quorum witness cannot be added to either iSCSI or non-Spectrum Accelerate target.

Deleting a remote target

Use the **target_delete** command to delete the definition of the specified remote target.

```
target_delete target=TargetName [ force_on_olvm_peer=<yes|no> ]
```

Parameters

Name	Type	Description	Mandatory	Default
target	Object name	Target that is deleted.	Y	N/A
force_on_olvm_peer	Boolean	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

Listing remote targets

Use the **target_list** command to list a specified remote target definition, or all target definitions.

```
target_list [ target=TargetName ] [ domain=DomainName ]
```

Parameters

Name	Type	Description	Mandatory	Default
target	Object name	Target name that is listed.	N	All targets
domain	Object name	The domain name.	N	All Domains

The following is listed for each target: port groups, ports, active/inactive status for each port, and the following mirroring-related values: max initialization rate, max resync rate, and max sync job rate.

Field ID	Field output	Default position
name	Name	1
scsi_type	SCSI Type	2
connected	Mirror Connectivity	3
ha_connected	HA Connectivity	4
max_initialization_rate	Max Initialization Rate	5
max_resync_rate	Max Resync Rate	6
max_syncjob_rate	Max Syncjob Rate	7
machine_serial_number	Target Serial Number	8
system_id	System ID	N/A
quorum_witness	Quorum Witness	9
xiv_target	XIV Target	N/A
iscsi_name	iSCSI Name	N/A
num_ports	Number of Ports	N/A
creator	Creator	N/A
connectivity_lost_event_threshold	Connection Threshold	N/A
peer_health	Peer Health	N/A
peer_health_reason	Peer Health Reason	N/A
peer_qw_configuration	Peer QW Configuration	N/A
coordinated_qw_lapse	Coordinated QW Lapse	N/A

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

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_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=Nextra2 fcaddress=10:00:00:17:38:27:ec:11
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed

User Category	Permission
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
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
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

Deleting a port from a remote system

Use the **target_port_delete** command to delete a port from the specified remote target.

```
target_port_delete target=TargetName < ipaddress=IPaddress | fcaddress=wwpn >
[ force_on_ha_peer=<yes|no> ]
```

Parameters

Name	Type	Description	Mandatory	
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	
force_on_ha_peer	Boolean	Force the deactivation on a HyperSwap target.	N	No

Example:

```
target_port_delete
target=Nextra2
fcaddress=10:00:00:17:38:27:ec:11
```

Output:

```
Command 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
- **TARGET_HAS_OLVM_RELATIONSHIP**

  Target has an IBM Hyper-Scale Mobility relationship - cannot be deactivate or deleted
- **TARGET_HAS_HA_RELATIONSHIP**

  Target has an IBM HyperSwap relationship - cannot be deactivate or deleted

Listing the ports of a remote target

Use the **target_port_list** command to list all ports of a target.

```
target_port_list [ target=TargetName ] [ domain=DomainName ]
```

Parameters

Name	Type	Description	Mandatory	Default
target	Object name	Target for which all ports should be listed.	N	All systems
domain	Object name	The domain name.	N	All Domains

Field ID	Field output	Default position
target_name	Target Name	1
scsi_type	Port Type	2
active	Active	3
fc_wwpn	WWPN	4
iscsi_ip_addr	iSCSI Address	5
iscsi_ip_port	iSCSI Port	6

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Renaming a remote target

Use the **target_rename** command to rename a remote target.

```
target_rename target=TargetName new_name=Name
```

Parameters

Name	Type	Description	Mandatory
target	Object name	Tne 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_512b_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 *system_id* variable on the remote system.	Y	
uses_512b_sectors	Boolean	Optimize the asynchronous mirror data transfer for remote targets with 512B sector size.	N	No

This command changes the system ID of the remote target.

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- **TARGET_BAD_NAME**

  Target name does not exist
- **TARGET_HAS_QUORUM_WITNESS_UPDATE_NOT_ALLOWED**

  Updating the target system_id is not allowed when there is a quorum witness defined on the target

Adding a Quorum Witness to a target

Use the **target_add_quorum_witness** command to attach a Quorum Witness to a remote target.

```
target_add_quorum_witness target=TargetName quorum_witness=QW_Name
```

Parameters

Name	Type	Description	Mandatory
target	Object name	The name of the target.	Y
quorum_witness	Object name	The name of the Quorum Witness that is associated with the target.	Y

Example:

```
target_add_quorum_witness target=t1 quorum=q1
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- **REMOTE_TARGET_SYSTEM_NOT_CONFIGURED_ON_QUORUM_WITNESS**

  The remote system is not configured on the quorum witness. Are you sure you want to configure the quorum witness '*Quorum Witness*' for target '*Target*?

  Troubleshooting: Configure the quorum witness on the remote system and try again.

Return codes

- **TARGET_BAD_NAME**

  Target name does not exist

- **QUORUM_WITNESS_BAD_NAME**

  Quorum witness name does not exist.

- **TARGET_HAS_A_QUORUM_WITNESS**

  target already has a quorum witness.

- **QUORUM_WITNESS_IS_NOT_ACTIVATED**

  The quorum witness is not activated.

- **QUORUM_WITNESS_CANNOT_BE_ADDED_TO_A_TARGET_OF_THIS_TYPE**

  A quorum witness cannot be added to either iSCSI or non-Spectrum Accelerate target.

Removing a Quorum Witness from a target

Use the **target_remove_quorum_witness** command to detach a Quorum Witness from a remote target.

```
target_remove_quorum_witness target=TargetName
```

Parameters

Name	Type	Description	Mandatory
target	Object name	The name of the target.	Y

Example:

```
target_remove_quorum_witness target=t1
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- **TARGET_BAD_NAME**

  Target name does not exist
- **TARGET_HAS_NO_QUORUM_WITNESS**

  The local target does not have a quorum witness defined.
- **TARGET_HAS_ENABLED_HA**

  There are HA relations with enabled HA configured with the target. Changing the quorum witness requires that HA high availability is disabled

Chapter 10. 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.

```
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

© Copyright IBM Corp. 2017

User Category	Permission	Condition
Storage integration administrator	Allowed	N/A
Application administrator	Conditionally Allowed	The volume is mapped to a host or a cluster associated with the user.
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
- **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 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_VOLUME_IS_SNAPSHOT**
  Slave volume is a snapshot
- **REMOTE_VOLUME_DATA_MIGRATION_UNSYNCHRONIZED**
  Data Migration has not completed to the remote volume
- **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
- **REMOTE_CONS_GROUP_EMPTY**
  Operation is not allowed on an empty Consistency Group(remote).
- **REMOTE_CONS_GROUP_MISMATCH**
  Remote Snapshot Group does not match Consistency Group volumes.
- **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_IS_OLVM_PROXY**
  The volume is in an IBM Hyper-Scale Mobility Proxy phase.

- **OPERATION_DENIED_OBJECT_MANAGED**

  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.
- **OPERATION_DENIED_REMOTE_OBJECT_MANAGED**

  The remote object is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.
- **CONS_GROUP_BAD_TARGET**

  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
- **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
- **MAX_SNAPSHOTS_PER_VOLUME_REACHED**

  Maximum number of snapshots already created
- **REMOTE_MAX_SNAPSHOTS_PER_VOLUME_REACHED**

  Maximum number of snapshots already created on remote
- **TARGET_SNAPSHOT_GROUP_BAD_NAME**

  Target Snapshot Group name does not exist.
- **REMOTE_DOMAIN_MAX_VOLUMES_REACHED**

  Maximum number of volumes already defined on remote machine domain
- **REMOTE_CONS_GROUP_BAD_NAME**

  Remote Consistency Group name does not exist.
- **SNAPSHOT_CAN_NOT_BE_CREATED_REMOTE_CONS_GROUP_IO_IS_NOT_PAUSED**

  The snapshot group will not be created since the Remote Consistency Group is not in a stopped state.
- **SNAPSHOT_CAN_NOT_BE_CREATED_REMOTE_CONS_GROUP_DEFINITION_CHANGED**

  The snapshot group will not be created since the volumes within the Remote Consistency Group have changed since the io_pause command was issued.
- **REMOTE_OVERWRITE_SNAPSHOT_GROUP_DOES_NOT_BELONG_TO_GIVEN_GROUP**

  Remote Snapshot Group belongs to another Consistency Group.

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_volume** or **cg_remove_volume** 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
- **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
- **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

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

User Category	Permission
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
- **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

Changing the designation of mirroring peers

Use the **mirror_change_designation** command to change the designation of mirroring peers: from primary to secondary, and vice versa.

```
mirror_change_designation < vol=VolName | cg=cgName > [ target=TargetName ]
 [ new_designation=<Primary|Secondary|None> ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	Master volume name.	N	N/A
cg	Object name	Master consistency group name.	N	N/A
target	Object name	Target mirror name. Mandatory if 2 mirrors are defined on the volume.	N	[none]
new_designation	Enumeration	The new designation of the peer If not specified, the command swaps the designation of the primary and secondary peer.	N	none

The command is issued on the master peer and affects both peers. The coupling has to be operational.

The designation change implied by this command reflects a decision to reset the designation of the mirroring peers, in contrast with the operational role, which is denoted by the master/dlave title.

There is no obligation to issue the command with a specification of the new designation. If the new designation is not specified, the command swaps the designations of both peers from their current value. The primary changes to secondary, and the secondary - to primary.

Example:

```
mirror_change_designation cg=reggie13_cg new_designation=Secondary
```

Output:

```
Command 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
- **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

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

- **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

Changing the role of a mirrored volume

Use the **mirror_change_role** command to change the role of a local mirroring peer.

```
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.

  - A warning is displayed: `Are you sure to change master to slave?"`
  - An event is generated
  - The master ceases accepting host requests
  - Completion of process is recorded in the log
- When applied on the slave, the slave becomes the master, starts accepting requests from hosts, and upon explicit activation starts replicating to the other peer (the original master).

If the synchronous mirroring is interrupted in the middle of the re-synchronization process, the slave volume may very probably be inconsistent. The last consistent image of the slave volume is preserved in the **last_consistent** snapshot (LCS), which is automatically created immediately before the re-synchronization starts. If the LCS exists, the command emits the warning *'Are you sure you want the mirror's local peer to become Master? The local peer has a last-consistent snapshot'*. In this case, the administrator must choose whether to use the existing contents of the previous slave volume, which may be inconsistent, or revert the previous slave volume to its **last_consistent** snapshot before issuing the **mirror_change_role** command.

Changing the roles in asynchronous mirroring:

- When 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

User Category	Permission
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_SYNCHED_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_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

- **SYSTEM_OUT_OF_PHYSICAL_SPACE**

  Operation not allowed while the system is out of physical space.

Changing a mirroring schedule for local peers

Use the **mirror_change_schedule** command to change the replication schedule for peers on the local system.

```
mirror_change_schedule < vol=VolName | cg=cgName > [ target=TargetName ] schedule=Schedule
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	Volume name on the local system.	N	N/A
cg	Object name	Consistency group name on the local system.	N	N/A
target	Object name	Target mirror name. Mandatory if 2 mirrors are defined on the volume.	N	[none]
schedule	Object name	A reference to a mirroring schedule	Y	N/A

This command changes the replication schedule for a peer on the local system. The new scheduling will become effective only if the peer is set as master.

Prerequisites:

- The coupling must be ASYNC_INTERVAL.
- The schedule's interval has to be shorter than the corresponding mirror's RPO.

The command fails under the following conditions:

- The specified target does not exist
- The specified target is non-mirrored
- The specified target is a volume that belongs to a mirrored consistency group
- The specified target synchronization type is not ASYNC_INTERVAL

Setting a scheduling reference:

- 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
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]

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 system does not copy the entire contents of the master volume upon activation. Instead, the system compares the contents of the master and the slave, and copies only the data, that is found to be different. This option allows you to reduce the time of initial synchronization when a huge amount of data is to be mirrored or when there is not enough data transfer bandwidth between the master and the slave system. For example, make tape backup copies of master volumes, restore these onto slave volumes, and create mirror relationships along with this option. The `offline` option does not require the slave volume to be formatted as opposed to the `online` option.

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

- **`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
- **NO_ASYNC_IN_THIN_PROVISIONED_POOL**
  Thin provisioned Pool cannot contain Volumes with Asynchronous Mirroring
- **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
- **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
- **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!
- **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.
- **REMOTE_VOLUME_HAS_MIRRORING_SNAPSHOTS**

  Remote Volume has snapshots created by previous mirroring process.
- **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
- **SYSTEM_OUT_OF_PHYSICAL_SPACE**

  Operation not allowed while the system is out of physical space.

Deactivating mirroring

Use the **mirror_deactivate** command to deactivate mirroring for a defined mirror coupling.

```
mirror_deactivate  < vol=<vol1[,vol2]...> |
    cg=cgName > [ target=TargetName ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	Master volume name or a list of master volumes.	N	N/A
cg	Object name	Master consistency group name or a list of master consistency groups.	N	N/A
target	Object name	Target mirror name. Mandatory if 2 mirrors are defined on the volume.	N	[none]

This command deactivates a coupling and switches it to the Inactive state. While in the Inactive state, only the master volume is updated, as opposed to the Active state, where the slave volume is updated together with the master volume.

The command cannot be issued on a slave.

If the mirroring is already inactive, this command has no effect and a success code is returned.

If more than one volume is specified, mirroring on all the volumes is deactivated. Furthermore, the deactivation of all the volumes is performed as an atomic operation, so that the slave volumes remain consistent with each other.

Deactivating a consistency group affects all of its volumes.

The command fails under the following conditions:

- The specified target does not exist.
- The specified target is non-mirrored.
- The specified target is a volume that belongs to a consistency group (in this case, the entire consistency group must be deactivated).
- Some of the specified targets are masters and some are slaves.
  - Each instance of the command can be applied to either master(s) or slave(s), but not to both.
- The target is a slave, yet the link is up.
- If multiple volumes are specified in the command and some are already part of an inactive mirror, the command will fail for all mirrors, including those that were active. The relevant return code is: **SYNC_ALREADY_INACTIVE**.

Access control

User Category	Permission
Storage administrator	Allowed
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
- **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

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
- **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
- **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

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)

Name	Type	Description	Mandatory	Default
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 in presented in yyyy-mm-dd hh:mm:ss format
- **Specified RPO**: the value in 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
crash_consistent	Crash Consistency	N/A	N/A
validate	Validation	N/A	N/A
is_standby	Standby	N/A	8

Output:

```
<command id="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="100777">
                <id value="100777"/>
                <creator value=""/>
                <creator_category value="none"/>
                <local_peer_id value="100776"/>
                <local_peer_name value="SYNC_vol_5"/>
                <schedule_name value=""/>
                <designation value="Secondary"/>
                <current_role value="Slave"/>
                <remote_mirror_id value="100872"/>
                <remote_peer_name value="SYNC_vol_4"/>
                <target_id value="100707"/>
                <target_name value="SYNC_target_2"/>
                <sync_type value="sync_best_effort"/>
                <sync_state value="Consistent"/>
                <active value="yes"/>
                <connected value="yes"/>
                <operational value="yes"/>
                <sync_progress value="100"/>
                <size_to_synchronize value="-1"/>
                <estimated_sync_time value="0"/>
                <mirror_error value="No_Error"/>
                <mirror_object value="Volume"/>
                <specified_rpo value=""/>
                <remote_rpo value=""/>
                <last_replicated_snapshot_time value=""/>
                <init_type value="online"/>
            </mirror>
        </return>
    </command>
</administrator>
<aserver status="DELIVERY_SUCCESSFUL"/>
</command>
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Obtaining statistics on past sync jobs

Use the **mirror_statistics_get** command to present statistics that are automatically gathered by the system on past sync jobs per specified mirrored volume or consistency job.

```
mirror_statistics_get <vol=VolName | cg=cgName> [ target=TargetName ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	Local volume name.	N	N/A
cg	Object name	Local consistency group name.	N	N/A
target	Object name	Target mirror name. Mandatory if 2 mirrors are defined on the volume.	N	[none]

The command output includes:

- Date and time created
- Date and time started to run
- Date and time finished
- Job size (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="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.
- **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
- **SYSTEM_OUT_OF_PHYSICAL_SPACE**
  Operation not allowed while the system is out of physical space.

Retrieving RPO thresholds

Use the **rpo_thresholds_get** command to list system RPO-related thresholds, that, once crossed, trigger the creation of a corresponding event.

```
rpo_thresholds_get
```

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:

- Only the following values are allowed in a schedule: 00:00:20, 00:00:30, 00:00:40,00:00:50, 00:01, 00:01:10, 00:01:20, 00:01:30, 00:01:40, 00:01:50, 00:02, 00:05, 00:10.
- A predefined schedule cannot be changed.

Outcome:

- If the update command is issued on a schedule that is not referenced by any object, a confirmation message is displayed.
- If the update command is issued on a schedule that is referenced to by an object (for example, mirroring couplings), a warning message is displayed.
- Sync jobs that are running will not be affected.
- Future sync jobs are scheduled based on the new schedule settings.

Example:

```
schedule_change interval=00:01 schedule=1min domain=* -y
```

Output:

```
Command 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 differnt 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.
- **ILLEGAL_TIME_SLOT_SPACE**

  Values should be between minimum to interval time

Creating a schedule object

Use the **schedule_create** command to define a schedule for replication.

```
schedule_create schedule=Schedule [ interval=IntervalSize ]
 [ type=<manual|interval|max|time> ] [ domain=DomainList ]
```

Parameters

Name	Type	Description	Mandatory	Default
schedule	Object name	The name of the schedule	Y	N/A
interval	N/A	The interval for asynchronous mirroring. Format: hh:mm [:ss].	N	00:10[:00]
type	Enumeration	The schedule type for asynchronous mirroring. Can be **manual** or **interval**.	N	interval
domain	N/A	The schedule will be attached to the specified domains. To specify several domains, separate them with a comma. To specify all existing domains, use "*".	N	none

This command creates a schedule definition. Schedules can be referenced to when specifying asynchronous mirroring couplings.

Limitations:

- Only the following values are allowed in a schedule: 00:00:20, 00:00:30, 00:00:40,00:00:50, 00:01, 00:01:10, 00:01:20, 00:01:30, 00:01:40, 00:01:50, 00:02, 00:05, 00:10.
- The system features a predefined schedule object with a non-user-configurable interval of 1 minute, named **min_interval**.

The **type** parameter:

Prior to the introduction of this parameter, each asynchronous mirror could be configured with an automatic schedule, whose interval specified how often a replication point and the corresponding replication process (sync job) should be automatically created. It was also possible to instruct the system to create a manual replication point and a corresponding sync job for a mirror using the dedicated CLI command **mirror_create_snapshot**. Finally, a single predefined schedule named *Never* with no interval settings was provided for mirrors that only required manual sync job creation.

The **type** parameter enables you to define multiple custom, user-configurable manual schedules. The creation of consistent, identical replication points for all mirrors set with such schedule, as well as corresponding sync jobs can be triggered using the dedicated CLI command **schedule_create_tick**, that specifies the schedule name as an argument. This facilitates external/scripted replication control for mirrors sharing the same schedule, without requiring them to be interval-based.

When **type=interval**, synchronization jobs for a mirror associated with the schedule will be triggered automatically, based on the specified interval.

When **type=manual**, synchronization jobs for a mirror associated with the schedule can be triggered by the command **schedule_create_tick**.

Once set, the schedule type cannot be changed.

Example:

```
schedule_create interval=00:01 schedule=1min domain=*
```

Output:

```
Command 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 defined only is exclude_time is set
- **ONLY_INTERVAL_SCHEDULE_MAY_HAVE_EXCLUSIONS**
  Exclusion period may defined only for Interval Schedule
- **DOMAIN_DOESNT_EXIST**
  Domain does not exist.
- **ILLEGAL_TIME_SLOT_SPACE**
  Values should be between minimum to interval time

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

Chapter 11. HyperSwap commands

This chapter describes the command line interface (CLI) for HyperSwap.

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

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

When certain conditions apply, an automatic and completely transparent failover is performed, so that the applications experience no downtime. As soon as the actual failure is recovered, the pair is automatically resynchronized.

As in other high availability solutions, HyperSwap requires a quorum witness component, to avoid split-brain situations. HyperSwap Quorum Witness is constantly monitoring the status of the related storage systems, and, if necessary, acts as a tiebreaker for conflict resolution.

The HyperSwap solution relies on Asymmetrical Logical Unit Access (ALUA) support to inform the host about the optimized paths to the storage system, and minimize I/O latency.

FlashSystem A9000 and FlashSystem A9000R HyperSwap capability does not require additional special hardware or software, and does not require any additional licensing.

Important: Some terminology used in this section is inconsistent with the terminology in other IBM FlashSystem A9000 and A9000R documentation and in IBM Hyper-Scale Module. For legacy purposes, this section still refers to "Master" and "Slave" volumes, which are elsewhere referred to as "Primary" and "Secondary". The new terminology is more suitable to describe the common replication technology, for both high availability (HyperSwap) and disaster recovery (Synchronous and Asynchronous mirroring).

Creating a HyperSwap relationship

Use the **ha_create** command to create a HyperSwap relationship.

```
ha_create < vol=VolName [ create_slave=<yes|no> [ remote_pool=RemotePoolName ] ]
 [ init_type=<online|offline> ] > | <cg=cgName slave_cg=SlaveCgName> target=TargetName
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	Local volume to be replicated (the Master).	N	N/A

© Copyright IBM Corp. 2017

Name	Type	Description	Mandatory	Default
create_slave	Boolean	Determines whether to create a new Slave volume or to use an existing one. If an existing remote volume is used, its name must match the local volume name.	N	no
remote_pool	Object name	The storage pool on the remote system. Relevant only if creating a Slave volume.	N	N/A
cg	Object name	The local consistency group to be mirrored.	N	N/A
slave_cg	Object name	The name of the Slave consistency group on the remote storage system.	N	N/A
target	Object name	The remote target to contain the Slave volume.	Y	N/A
init_type	Enumeration	The initialization method of the Slave volume.	N	online

This command defines a new HyperSwap relationship between Master and Slave peers.

When you define a HyperSwap relationship, the following rules are enforced:

- The Slave volume is not mapped
- The Master and Slave peers must be configured with the same QW
- The Master and Slave connectivity to the QW must be healthy
- If the HyperSwap relationship includes an existing Slave volume, the Slave volume's name must match the Master volume's name

As part of the operation, Master volume metadata is copied to the Slave volume, and the Slave volume identity changes.

A HyperSwap relationship is created in the Standby state. It must then be activated in order to start the initialization process, which copies data from the Master to the Slave.

The following initialization methods are available:

- The `online` option (default) enables an over-the-wire initialization. In other words, it uses an inter-site link to replicate the Master's initial state to the Slave, starting once HyperSwap is first activated (**ha_activate**). During initialization, the HyperSwap relationship status will be *Initializing*.
- If the `offline` option is selected, the initialization of the Slave peer is not done by replicating the Master's initial image, but rather by creating its offline replica. In other words, it restores to the Slave a mirror image that is backed up on the Master. Once the relationship is activated, the contents of the volumes are

compared, and only modified data is synchronized over the wire. This process is usually much faster than online initialization.

During initialization, the HyperSwap relationship status will be *Initializing*.

A storage system can have multiple HyperSwap relationships between pairs of peers on various remote systems. However, when the peers have consistency group HyperSwap relationships, all the volumes included in a specific consistency group HyperSwap relationship can only be replicated between one pair of storage systems. Therefore, when a volume peer on a storage system (for example: A) has a HyperSwap 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 HyperSwap relationship with a volume on storage system B. The same is true for volumes from storage system B to A. In addition, the HyperSwap consistency group has one sync job for all pertinent HyperSwap volumes within the consistency group.

Prior to issuing this command on a consistency group, make sure that the consistency group is empty.

Example:

```
ha_create vol=regular_volume target=123 create_slave=yes remote_pool=333
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- **VOLUME_SIZE_VERY_LARGE_ARE_YOU_SURE**

  Volume size is very large. It may not be possible to mirror this volume to older versions of XIV. Are you sure?

Return codes

- **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.
- **LOCAL_MAX_HA_REACHED**
  Maximum number of HA already defined on local machine
- **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
- **VOLUME_HAS_OLVM**
  IBM Hyper-Scale Mobility relation is defined for this volume
- **VOLUME_HAS_HA**
  This operation is forbidden on a volume with HA relation.
- **TARGET_VOLUME_HAS_OLVM**
  This target volume is part of an IBM Hyper-Scale Mobility relation.
- **TARGET_VOLUME_HAS_HA**
  This operation is forbidden with a target volume with HA relation.
- **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_HA_REACHED**
  Maximum number of HA 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
- **REMOTE_CONS_GROUP_IS_MIRRORED**
  Remote Consistency Group has mirroring defined for it.
- **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
- **HA_INCOMPATIBLE_TARGET_VERSION**
  HA 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.
- **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
- **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

- **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!
- **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.
- **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
- **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.
- **REMOTE_VOLUME_HAS_MIRRORING_SNAPSHOTS**
  Remote Volume has snapshots created by previous mirroring process.
- **DATA_REDUCTION_TIER_IS_OFFLINE**
  Data reduced tier is offline, operation is not allowed
  Troubleshooting: Contact IBM Support
- **SYSTEM_OUT_OF_PHYSICAL_SPACE**
  Operation not allowed while the system is out of physical space.
- **REMOTE_DATA_REDUCTION_TIER_IS_OFFLINE**
  Data reduced tier of remote system is offline, operation is not allowed
  Troubleshooting: Contact IBM Support
- **REMOTE_SYSTEM_OUT_OF_PHYSICAL_SPACE**
  Operation not allowed while the remote system is out of physical space.
- **SLAVE_VOLUME_IS_MAPPED**
  The Slave volume is mapped
- **TARGET_HAS_NO_QUORUM_WITNESS**
  The local target does not have a quorum witness defined.
- **REMOTE_TARGET_HAS_NO_QUORUM_WITNESS**
  The target on the remote system does not have a quorum witness defined.

- **HA_TARGET_QUORUM_WITNESS_IS_NOT_ACTIVATED**
  The quorum witness associated with the target is deactivated.
- **HA_REMOTE_TARGET_QUORUM_WITNESS_IS_NOT_ACTIVATED**
  The quorum witness associated with the remote target is deactivated.
- **HA_CONNECTIVITY_NOT_SUFFICIENT**
  The connectivity between the systems is not sufficient for HA.
- **HA_LOCAL_PEER_HAS_NO_QUORUM_WITNESS_CONNECTIVITY**
  The local peer connectivity to the QW is not operational.
- **HA_REMOTE_PEER_HAS_NO_QUORUM_WITNESS_CONNECTIVITY**
  The remote peer connectivity to the QW is not operational.
- **MAX_NUM_OF_PROXY_VOLUME_REACHED**
  Failed to create mirror, max number of proxy volumes exceeded.
- **REMOTE_MAX_NUM_OF_PROXY_VOLUME_REACHED**
  Failed to create mirror, remote max number of proxy volumes exceeded.
- **REMOTE_MAX_METADATA_OBJECTS_REACHED**
  Maximal number of metadata objects has been reached on a remote system.
- **VOLUME_BELONGS_TO_MIRRORED_CONS_GROUP**
  Volume mirror is part of Consistency Group mirror.
- **HA_PEER_QUORUM_WITNESS_CONFIGURATION_NOT_VERIFIED**
  The peer QW connectivity is not verified
- **HA_REMOTE_PEER_QUORUM_WITNESS_CONFIGURATION_NOT_VERIFIED**
  The remote peer QW connectivity is not verified
- **TARGET_PEER_NOT_HEALTHY**
  Target peer health not marked as healthy
  Troubleshooting: check quorum witness configuration
- **HOST_TYPE_IS_NOT_CONFIGURED**
  Cannot associate a hyperswap volume with a host with unconfigured type. IMPORTANT: Please read the Hyperswap chapter in the 'Best Practice' document to understand the solution requirements.

Viewing the status of HyperSwap volumes and consistency groups

Use the **ha_list** command to display the status of HyperSwap volumes and consistency groups.

```
ha_list [ < [ vol=VolName ] [ target=TargetName ] > | cg=cgName | scope=<cg|volume> ]
 [ domain=DomainName ]
```

Parameters

Name	Type	Description	Mandatory	Default
scope	Enumeration	List type: all HyperSwap volumes and consistency groups.	N	All (if no value is specified)
vol	Object name	Local volume name.	N	[none]

Name	Type	Description	Mandatory	Default
cg	Object name	Local consistency group name.	N	[none]
target	Object name	Remote target name.	N	[none]
domain	Object name	The domain name.	N	All Domains

The size, part, and time to synchronize are unknown if this is the Slave and connection is broken.

The following default parameters are shown:

- **Name**
- **HA Object**: Volume or CG
- **Role**: Master or Slave
- **Remote System**: The name of the target storage system
- **Active**: Yes or No
- **Status**: Initializing, Synchronized, or Unsynchronized
- **Link Up**: Yes or No
- **Automatic Failover**: The ability of a Slave volume to perform automatic failover: Active, Inactive, or N/A. N/A is returned for a Master volume or consistency group, or for a volume that is part of a consistency group.

The following optional parameters can be listed by explicitly specifying the proper columns:

- **Designation**: Primary or Secondary
- **Estimated Sync Time**: Estimated time (in seconds) for synchronization to complete. This parameter is not available for a Slave volume or consistency group, if its **Link Up** status is No.
- **Size To Synchronize**: The amount of data (in MB) to synchronize. This parameter is not available for a Slave volume or consistency group, if its **Link Up** status is No.
- **Operational**: Yes or No
- **Synch Progress**: The amount of synchronized data (in %)
- **Automatic Failover Reason** indicates why an automatic failover has been performed: N/A, User setting, part of a CG, Unsynced, Quorum Witness issue
- **IO Service** indicates the local peer ability to serve I/O: Active or Unavailable

Example:

```
ha_list
```

Output:

```
Name   HA Object   Role     Remote System         Active   Status           Link Up
------ ----------- -------- --------------------- -------- ---------------- ---------
ha_1   Volume      Master   target-3726085-0008   yes      Synchronized     yes
ha_2   Volume      Slave    target-3726085-0008   yes      Synchronized     yes
ha_3   Volume      Slave    target-3726085-0008   yes      Initializing     yes

Cont.:

Automatic Failover
--------------------
N/A
Active
Inactive
```

Field ID	Field output	Description	Default position
local_peer_name	Name	N/A	1
ha_object	HA Object	N/A	2
current_role	Role	N/A	3
target_name	Remote System	N/A	4
active	Active	N/A	5
sync_state	Status	N/A	6
ha_connected	Link Up	N/A	7
automatic_failover	Automatic Failover	N/A	8
automatic_failover_reason	Automatic Failover Reason	N/A	N/A
remote_peer_name	Remote Peer	N/A	N/A
designation	Designation	N/A	N/A
size_to_synchronize	Size To Sync (MiB)	N/A	N/A
operational	Operational	N/A	N/A
sync_progress	Sync Progress (%)	N/A	N/A
mirror_error	Mirror Error	No Error, Secondary pool exhausted, Configuration error or No thin provisioning resources	N/A
crash_consistent	Crash Consistency	N/A	N/A
validate	Validation	N/A	N/A
ha_high_availability_state	HA High Availability State	N/A	N/A
ha_unavailable_reason	HA Unavailable Reason	N/A	N/A
ha_sync_state	HA Sync State	N/A	N/A
ha_object_state	HA object State	N/A	N/A
io_service	I/O Service	N/A	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed

User Category	Permission
Technicians	Disallowed

Activating a HyperSwap relationship

Use the **ha_activate** command to activate a HyperSwap relationship of volumes or consistency groups.

```
ha_activate < vol=VolName | cg=cgName > [ target=TargetName ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	Master volume name.	N	N/A (either a volume or consistency must be specified)
cg	Object name	Master consistency group name.	N	N/A (either a volume or consistency must be specified)
target	Object name	Target HyperSwap relationship name.	N	[none]

The command updates the Quorum Witness about the state of the HyperSwap relationship after the operation. If HyperSwap is enabled, the command will fail if the Quorum Witness update fails. If HyperSwap is disabled, the command may succeed even if the Quorum Witness update fails.

The following is required for a successful command completion:

- The Master and Slave are configured on the same Quorum Witness
- The connectivity of Master and Slave to the Quorum Witness is healthy
- The specified volume is a Master
- The connectivity between Master and Slave is sufficient for HyperSwap (that is, the **link up** status in the **ha_list** command output is *yes*).

If the relationship is already in the Active state, nothing is done and a success code is returned.

A HyperSwap relationship cannot be activated, if:

- The command is issued on a Master that did not receive acknowledgment from the Slave following the **cg_add_volume** or **cg_remove_volume** command, due to the command's timeout or to an unexpected failure. In this case, the command fails and the **HA_CONFIGURATION_ERROR** code is returned. This means that the member lists of the HyperSwap consistency group peers are not the same.
- The command is issued on a Master that did not receive acknowledgment from the Slave following a **vol_resize** command, due to the command's timeout or to an unexpected failure. In this case, the command fails and the **HA_CONFIGURATION_ERROR** code is returned. This error means that the sizes of the HyperSwap volume peers are not the same.

Example:

```
ha_activate vol=regular_volume target=123
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- **VOLUME_BAD_NAME**

  Volume name does not exist
- **CONS_GROUP_BAD_NAME**

  Consistency Group name does not exist.
- **LOCAL_PEER_IS_NOT_MASTER**

  Local peer is not the master
- **HA_CONFIGURATION_ERROR**

  HA local configuration does not match remote configuration
- **REMOTE_MAX_VOLUMES_REACHED**

  Maximum number of volumes already defined on remote machine
- **SYNC_ALREADY_ACTIVE**

  Synchronization is already active
- **VOLUME_BELONGS_TO_HA_CONS_GROUP**

  This volume belongs to a Ha consistency group. The operation can be carried out on a volume which is not part of a consistency group, or on the consistency group itself.
- **DOMAIN_MAX_VOLUMES_REACHED**

  The domain exceeds the maximum allowed number of volumes.
- **HA_CAN_NOT_BE_ACTIVATED**

  HA cannot be activated
- **HA_CONS_GROUP_MEMBERSHIP_MISMATCH**

  HA 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.
- **HA_SIZE_MISMATCH**

  HA Slave volume and HA Master Volume sizes are different
- **HA_RETRY_OPERATION**

There is an operation in progress on this HA , 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_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_DOMAIN_MAX_VOLUMES_REACHED**

  Maximum number of volumes already defined on remote machine domain
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  Data reduced tier is offline, operation is not allowed

  Troubleshooting: Contact IBM Support
- **REMOTE_DATA_REDUCTION_TIER_IS_OFFLINE**

  Data reduced tier of remote system is offline, operation is not allowed

  Troubleshooting: Contact IBM Support
- **REMOTE_SYSTEM_OUT_OF_PHYSICAL_SPACE**

  Operation not allowed while the remote system is out of physical space.
- **TARGET_PEER_NOT_HEALTHY**

  Target peer health not marked as healthy

  Troubleshooting: check quorum witness configuration
- **TARGET_HAS_NO_QUORUM_WITNESS**

  The local target does not have a quorum witness defined.
- **REMOTE_TARGET_HAS_NO_QUORUM_WITNESS**

  The target on the remote system does not have a quorum witness defined.
- **VOLUME_IS_NOT_HA**

  This is not a high availability volume.
- **CONS_GROUP_IS_NOT_HA**

  Local Consistency Group does not have HA definitions.
- **HA_CONNECTIVITY_NOT_SUFFICIENT**

  The connectivity between the systems is not sufficient for HA.
- **HA_PEER_QUORUM_WITNESS_CONFIGURATION_NOT_VERIFIED**

  The peer QW connectivity is not verified
- **HA_REMOTE_PEER_QUORUM_WITNESS_CONFIGURATION_NOT_VERIFIED**

  The remote peer QW connectivity is not verified
- **HA_TARGET_QUORUM_WITNESS_IS_NOT_ACTIVATED**

  The quorum witness associated with the target is deactivated.
- **HA_REMOTE_TARGET_QUORUM_WITNESS_IS_NOT_ACTIVATED**

  The quorum witness associated with the remote target is deactivated.

Deactivating a HyperSwap relationship

Use the **ha_deactivate** command to deactivate HyperSwap volumes or consistency groups.

```
ha_deactivate  < vol=<vol1[,vol2]...> | cg=cgName > [ target=TargetName ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	Master volume name or a list of master volumes.	N	N/A (either a volume or consistency must be specified)
cg	Object name	Master consistency group name or a list of master consistency groups.	N	N/A (either a volume or consistency must be specified)
target	Object name	Target HyperSwap relationship name.	N	[none]

This command deactivates a HyperSwap relationship and changes its status to *Inactive*. While in the *Inactive* state, only the Master volume is updated, as opposed to the *Active* state, where the Slave volume is updated together with the Master volume.

Example:

```
ha_deactivate vol=regular_volume -y
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- **VOLUME_BAD_NAME**

  Volume name does not exist
- **VOLUME_IS_NOT_HA**

  This is not a high availability volume.
- **CONS_GROUP_BAD_NAME**

Consistency Group name does not exist.

- **CONS_GROUP_IS_NOT_HA**

  Local Consistency Group does not have HA definitions.
- **LOCAL_PEER_IS_NOT_MASTER**

  Local peer is not the master
- **SYNC_ALREADY_INACTIVE**

  Synchronization is already inactive
- **VOLUME_BELONGS_TO_HA_CONS_GROUP**

  This volume belongs to a Ha consistency group. The operation can be carried out on a volume which is not part of a consistency group, or on the consistency group itself.
- **HA_RETRY_OPERATION**

  There is an operation in progress on this HA , 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_TARGET_MISMATCH**

  Volume and target do not match
- **CONS_GROUP_BAD_TARGET**

  Target name does not match the Consistency Group
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  Data reduced tier is offline, operation is not allowed

  Troubleshooting: Contact IBM Support

Deleting a HyperSwap relationship

Use the **ha_delete** command to delete a HyperSwap relationship.

```
ha_delete < vol=VolName | cg=cgName > [ target=TargetName ] [ force_on_slave=<Yes|No> ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	Local Master volume name.	N	N/A (either a volume or consistency must be specified)
cg	Object name	Local Master consistency group name.	N	N/A (either a volume or consistency must be specified)
target	Object name	Target HyperSwap relationship name	N	[none]
force_on_slave	Boolean	Forces the deletion of the HyperSwap relationship on the Slave peer. Available only if the Slave is in the *Initializing* state.	N	no

The command can only be issued on the Master. It deletes only the definition of the HyperSwap relationship. Neither the volumes themselves, nor their snapshots are deleted.

To delete a HyperSwap relationship, the following preconditions must be met:

- The HyperSwap relationship is in the *Standby* state, in other words, it has just been created or deactivated.
- The communication is established. If there is no communication, the HyperSwap relationship is only deleted on the Master. When the communication resumes, a configuration error will be issued on the Slave.
- The Slave volume is not mapped.

After deleting a HyperSwap relationship, both its peers are labeled as *none*, meaning that they are no longer configured as either Master or Slave.

After the command's successful completion, the Slave volume's SCSI identification is replaced. The volume name, external ID, lock state, and metadata remain identical to the Master volume values.

Deletion when the HyperSwap relationship is inactive or when the connectivity has failed

On an active Master volume

On an active Master volume, the command can be executed as described above.

On an active Slave volume

Important: Never map the Slave volume to a host. If you need to read or update data, use the **vol_copy** command to copy the data to a new volume, and map this new volume to the host.

Deleting a HyperSwap relationship when the communication between the peers is down, deletes only the HyperSwap relationship from the Master. To delete the HyperSwap relationship from the Slave:

- Run the **ha_change_role** command to turn the Slave into the Master
- Run **ha_delete**.

The **force_on_slave** parameter can be used only if the HyperSwap relationship is in the *Initializing* phase. In any other state, change the Slave into the Master, and run **ha_delete**.

On a disconnected Master or Slave volume

To delete the HyperSwap relationship on a disconnected peer volume, re-connect the peer volume and run **ha_delete**.

If re-connecting the peer volume is not possible, contact IBM support.

Example:

```
ha_delete vol=regular_volume -y
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_DELETE_CG_HA**

  Are you sure you want to delete the HA relationships of the CG and of all volumes in the CG?

Return codes

- **VOLUME_BAD_NAME**

  Volume name does not exist
- **VOLUME_IS_NOT_HA**

  This is not a high availability volume.
- **CONS_GROUP_IS_NOT_HA**

  Local Consistency Group does not have HA definitions.
- **CONS_GROUP_BAD_NAME**

  Consistency Group name does not exist.
- **LOCAL_PEER_IS_NOT_MASTER**

  Local peer is not the master
- **HA_IS_ACTIVE**

  HA is currently active
- **FORCE_DELETE_NOT_ALLOWED_ON_MASTER**

  Only slave mirrors need to be forced to be deleted
- **VOLUME_BELONGS_TO_HA_CONS_GROUP**

  This volume belongs to a Ha consistency group. The operation can be carried out on a volume which is not part of a consistency group, or on the consistency group itself.
- **HA_RETRY_OPERATION**

  There is an operation in progress on this HA , please try again your request in a few seconds

  Troubleshooting: Please try again the command in a few seconds
- **HA_IS_NOT_INITIALIZING**

  Operation ha_delete is permitted only during the Initialization phase.
- **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
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  Data reduced tier is offline, operation is not allowed

  Troubleshooting: Contact IBM Support
- **SLAVE_VOLUME_IS_MAPPED**

  The Slave volume is mapped
- **VOLUME_IS_AN_UNAVAILABLE_HYPERSWAP_PEER**

  The operation is not permitted on a HyperSwap target which is unavailable for I/O.

Switching roles between Master and Slave volumes

Use the **ha_switch_roles** command to switch roles between Master and Slave volumes.

```
ha_switch_roles <vol=VolName | cg=cgName> [ target=TargetName ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	Local volume name.	N	N/A (either a volume or consistency must be specified)
cg	Object name	Local consistency group name.	N	N/A (either a volume or consistency must be specified)
target	Object name	N/A	N	[none]

The command can only be issued on the Master volume, and only if the HyperSwap relationship is activated and synchronized.

When this command is issued, the system performs all pending writes, and only after all pending writes have been committed, the roles are switched.

Following the successful completion of the command:

- The volume that was previously the Master becomes the Slave
- The volume that was previously the Slave becomes the Master
- The HyperSwap relationship remains active.

Example:

```
ha_switch_roles <vol=VolName | cg=cgName> [ target=TargetName ]
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_SWITCH_THE_PEER_ROLES**

  Switching the roles may lead to inactive automatic failover as a result of the intended slave failed connectivity with the Quorum Witness. Are you sure you want to continue?

Return codes

- **VOLUME_BAD_NAME**

  Volume name does not exist
- **CONS_GROUP_BAD_NAME**

  Consistency Group name does not exist.
- **LOCAL_PEER_IS_NOT_MASTER**

  Local peer is not the master
- **HA_IS_NOT_SYNCHRONIZED**

  HA is not synchronized
- **VOLUME_HAS_DATA_MIGRATION**

  Data Migration is defined for this volume
- **REMOTE_TARGET_NOT_CONNECTED**

  There is currently no connection from the target system
- **VOLUME_BELONGS_TO_HA_CONS_GROUP**

  This volume belongs to a Ha consistency group. The operation can be carried out on a volume which is not part of a consistency group, or on the consistency group itself.
- **HA_CONNECTIVITY_NOT_SUFFICIENT**

  The connectivity between the systems is not sufficient for HA.
- **HA_RELATION_MASTER_COULD_NOT_UPDATE_QW_AFTER_RETURN_TO_GOOD_STATE**

  The Master had control of the relation during a past failure, and was unable to update the quorum witness after the failure was recovered
- **HA_HAS_SYNC_JOB**

  Operation is not permitted on a HA with active sync jobs
- **HA_RETRY_OPERATION**

  There is an operation in progress on this HA , please try again your request in a few seconds

  Troubleshooting: Please try again the command in a few seconds
- **REMOTE_HA_IS_NOT_ACTIVE**

  Remote HA is not active
- **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
- **COMMAND_NOT_SUPPORTED_FOR_OLVM_VOLUMES**

  This command is not supported for IBM Hyper-Scale Mobility volumes.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  Data reduced tier is offline, operation is not allowed

  Troubleshooting: Contact IBM Support
- **SYSTEM_OUT_OF_PHYSICAL_SPACE**

  Operation not allowed while the system is out of physical space.
- **VOLUME_IS_NOT_HA**

  This is not a high availability volume.
- **CONS_GROUP_IS_NOT_HA**

  Local Consistency Group does not have HA definitions.

Changing a peer role in a HyperSwap volume

Use the **ha_change_role** command to change the role of a local HyperSwap relationship peer from Master to Slave or from Slave to Master.

```
ha_change_role <vol=VolName | cg=cgName> [ target=TargetName ]
 [ new_role=<Master|Slave|None> ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	Local volume name. Must be specified if the command is applied to a volume.	N	N/A
cg	Object name	CG name. Must be specified if the command is applied to a consistency group.	N	N/A
target	Object name	Target HyperSwap relationship name.	N	[none]
new_role	Enumeration	Role name of the peer. If not specified, the command swaps peer roles between Master and Slave.	N	none

This command changes the role of the local peer from Master to Slave or from Slave to Master when the HyperSwap relationship is not activated. The command

should be issued on both peers before the relationship is activated again, so that upon reconnection there still will be one Master and one Slave.

This command is used during recovery after an automatic failover, or in order to perform a manual failover when the automatic failover did not take place.

To successfully change the Master to the Slave, the following conditions must be met:

- The activation state must be *Standby*
- The volume can be in any phase, except *Initializing*.

After the Master is successfully changed to the Slave:

- All changes made to the Master since the last peer synchronization are reverted to their original value. The Master ceases serving host requests, and is set to accept replication from the other peer as a Slave.
- An event is generated
- The Master ceases accepting host requests
- The process completion is recorded in the log

To successfully change the Slave to the Master, the following conditions must be met:

- The relationship is deactivated
- The relationship is synced, otherwise some data may be lost. If the relationship is not synced, a warning is issued, prompting you to confirm before proceeding.

Before changing the Slave to the Master:

- Make sure that the original Master is not available and cannot become available while the other peer is a Master. To verify this, run the **ha_list** command on the Master, and check the value of the attribute *IO Service*. Only if the returned value is *Unavailable*, proceed with issuing the **ha_change_role** command on the Slave.
- Stop the applications using the HyperSwap volume(s). Note that each application must be stopped completely and not merely paused in order to make sure that it does not use any cached state when accessing the volume for the first time after the role change.

Note: Failure to fulfill both of the above requirements may result in a data integrity issue.

After the Slave is successfully changed to the Master, the volume starts accepting requests from hosts. Upon explicit activation, it starts replicating to the other peer (the original Master).

If the synchronous mirroring is interrupted in the middle of the re-synchronization process, the Slave volume may very probably be inconsistent. The last consistent image of the Slave volume is preserved in the **last_consistent** snapshot (LCS), which is automatically created immediately before the re-synchronization starts. If the LCS exists, the command emits the warning *'Are you sure you want the mirror's local peer to become Master? The local peer has a last-consistent snapshot'*. In this case, the administrator must choose whether to use the existing contents of the previous Slave volume, which may be inconsistent, or revert the previous Slave volume to its **last_consistent** snapshot before issuing the **ha_change_role** command.

Example:

```
ha_change_role vol=regular_volume
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- **SOME_DATA_WILL_BE_LOST_ARE_YOU_SURE**

  Are you sure you want the mirror'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_IS_NOT_HA**

  This is not a high availability volume.
- **CONS_GROUP_IS_NOT_HA**

  Local Consistency Group does not have HA definitions.
- **HA_IS_ACTIVE**

  HA is currently active
- **HA_IS_INITIAL**

  Operation is not permitted during the HA Initialization phase.
- **VOLUME_BELONGS_TO_HA_CONS_GROUP**

  This volume belongs to a Ha consistency group. The operation can be carried out on a volume which is not part of a consistency group, or on the consistency group itself.
- **VOLUME_BAD_NAME**

  Volume name does not exist
- **CONS_GROUP_BAD_NAME**

  Consistency Group name does not exist.
- **VOLUME_HAS_DATA_MIGRATION**

  Data Migration is defined for this volume
- **HA_RETRY_OPERATION**

  There is an operation in progress on this HA , please try again your request in a few seconds

Troubleshooting: Please try again the command in a few seconds

- **HA_HAS_NO_SYNCHED_SNAPSHOT**
  HA does not have a synchronized Snapshot.
- **MASTER_CANNOT_BE_DEMOTED**
  Master cannot be demoted to Slave role, Peer status mismatch
- **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
- **COMMAND_NOT_SUPPORTED_FOR_OLVM_VOLUMES**
  This command is not supported for IBM Hyper-Scale Mobility volumes.
- **DATA_REDUCTION_TIER_IS_OFFLINE**
  Data reduced tier is offline, operation is not allowed
  Troubleshooting: Contact IBM Support
- **SYSTEM_OUT_OF_PHYSICAL_SPACE**
  Operation not allowed while the system is out of physical space.

Restoring the availability of a Master volume

Use the **ha_restore_availability** command to restore the availability of a Master volume, that became unavailable due to a failure.

```
ha_restore_availability <vol=VolName | cg=cgName>
```

Parameters

Name	Type	Description	Mandatory
vol	Object name	Local volume name. Must be specified if the command is applied to a volume.	N
cg	Object name	CG name Must be specified if the command is applied to a consistency group.	N

As a result of some failure scenarios, the Master may assume that the Slave has performed an automatic failover, and stop handling I/O. In such cases, the user can choose to disable the remote peer and manually restore the availability of the Master.

Prior to issuing this command, make sure that the remote target is a Slave. If it is a Master, make sure that it is not available. To verify this, run the **ha_list** command on the remote target, and check the value of the attribute *IO Service*. Only if the returned value is *Unavailable*, proceed with issuing the **ha_restore_availability** command.

Upon issuing this command, the following occurs:

- An event is generated
- The Master volume becomes available.

Example:

```
ha_restore_availability vol=regular_volume
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- **VOLUME_IS_NOT_HA**

  This is not a high availability volume.
- **CONS_GROUP_IS_NOT_HA**

  Local Consistency Group does not have HA definitions.
- **CONS_GROUP_BAD_NAME**

  Consistency Group name does not exist.
- **HA_IS_ACTIVE**

  HA is currently active
- **VOLUME_BELONGS_TO_HA_CONS_GROUP**

  This volume belongs to a Ha consistency group. The operation can be carried out on a volume which is not part of a consistency group, or on the consistency group itself.
- **VOLUME_BAD_NAME**

  Volume name does not exist
- **LOCAL_PEER_IS_NOT_MASTER**

  Local peer is not the master
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  Data reduced tier is offline, operation is not allowed

  Troubleshooting: Contact IBM Support
- **SYSTEM_OUT_OF_PHYSICAL_SPACE**

  Operation not allowed while the system is out of physical space.

Creating a HyperSwap volume snapshot (ad hoc sync job)

Use the **ha_create_snapshot** command to simultaneously create snapshots on both peers of a HyperSwap relationship.

```
ha_create_snapshot <vol=VolName | cg=cgName> [ target=TargetName ]
 name=Name  [ slave_name=SnapshotName ] [ delete_priority=del_value ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	The name of the volume to create a snapshot for.	N	N/A
cg	Object name	Local master consistency group name.	N	N/A
target	Object name	Target HyperSwap relationship name.	N	[none]
name	Object name	The name of the new snapshot.	Y	N/A
slave_name	Object name	The name of the new snapshot on the slave.	N	[none]
delete_priority	Integer	The deletion priority of the volume's snapshot.	N	1

This command takes a snapshot of the source peer (Master) and the target peer (Slave) at exactly the same time. The snapshots created concurrently on the Master and Slave are identical.

Pre-requisite:

- The HyperSwap relationship is activated and synchronized.

The snapshots created by this command can be managed with regular snapshot commands. For example, to delete these snapshots, issue the **snapshot_delete** command at each peer.

Example:

```
ha_create_snapshot <vol=VolName | cg=cgName> [ target=TargetName ]
 name=Name  [ slave_name=SnapshotName ] [ delete_priority=del_value ]
```

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	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.
- **VOLUME_IS_NOT_HA**

  This is not a high availability volume.
- **CONS_GROUP_IS_NOT_HA**

  Local Consistency Group does not have HA definitions.
- **LOCAL_PEER_IS_NOT_MASTER**

  Local peer is not the master
- **HA_IS_NOT_SYNCHRONIZED**

  HA is not synchronized
- **HA_RETRY_OPERATION**

  There is an operation in progress on this HA , please try again your request in a few seconds

  Troubleshooting: Please try again the command in a few seconds
- **HA_IS_NON_OPERATIONAL**

  Ha 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_VOLUME_IS_SNAPSHOT**

  Slave volume is a snapshot
- **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
- **REMOTE_CONS_GROUP_EMPTY**

  Operation is not allowed on an empty Consistency Group(remote).
- **REMOTE_CONS_GROUP_MISMATCH**

  Remote Snapshot Group does not match Consistency Group volumes.
- **SNAPSHOT_ILLEGAL_PRIORITY**

  Illegal snapshot priority; must be an integer between 1 and 4.
- **SNAPSHOT_IS_INTERNAL**

  Internal snapshots cannot be mapped, modified or deleted.
- **SNAPSHOT_GROUP_IS_INTERNAL**

  Internal snapshots cannot be mapped, modified 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
- **VOLUME_BAD_PREFIX**

  Volume name has a reserved prefix
- **VOLUME_BELONGS_TO_HA_CONS_GROUP**

  This volume belongs to a Ha consistency group. The operation can be carried out on a volume which is not part of a consistency group, or on the consistency group itself.
- **VOLUME_EXISTS**

Volume name already exists

- **VOLUME_BAD_NAME**

  Volume name does not exist
- **VOLUME_IS_NOT_CONSISTENT_SLAVE**

  Operation not allowed on slave volume that is not consistent.
- **VOLUME_IS_SNAPSHOT**

  Operation is not permitted on snapshots
- **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
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  Data reduced tier is offline, operation is not allowed

  Troubleshooting: Contact IBM Support
- **MAX_SNAPSHOTS_PER_VOLUME_REACHED**

  Maximum number of snapshots already created
- **REMOTE_MAX_SNAPSHOTS_PER_VOLUME_REACHED**

  Maximum number of snapshots already created on remote
- **REMOTE_DOMAIN_MAX_VOLUMES_REACHED**

  Maximum number of volumes already defined on remote machine domain
- **REMOTE_CONS_GROUP_BAD_NAME**

  Remote Consistency Group name does not exist.
- **SNAPSHOT_CAN_NOT_BE_CREATED_REMOTE_CONS_GROUP_IO_IS_NOT_PAUSED**

  The snapshot group will not be created since the Remote Consistency Group is not in a stopped state.
- **SNAPSHOT_CAN_NOT_BE_CREATED_REMOTE_CONS_GROUP_DEFINITION_CHANGED**

  The snapshot group will not be created since the volumes within the Remote Consistency Group have changed since the io_pause command was issued.
- **REMOTE_OVERWRITE_SNAPSHOT_GROUP_DOES_NOT_BELONG_TO_GIVEN_GROUP**

  Remote Snapshot Group belongs to another Consistency Group.

Changing the designation of HyperSwap relationship peers

Use the **ha_change_designation** command to change the designation of HyperSwap relationship peers from Primary to Secondary and vice versa.

```
ha_change_designation < vol=VolName | cg=cgName > [ target=TargetName ]
 [ new_designation=<Primary|Secondary|None> ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	Master volume name.	N	N/A
cg	Object name	Master consistency group name.	N	N/A

Name	Type	Description	Mandatory	Default
target	Object name	The name of the target.	N	[none]
new_designation	Enumeration	The new designation of the peer. If not specified, the command swaps the designation of the Primary and Secondary peers.	N	none

The designation in a HyperSwap relationship reflects the user's decision where the Primary (Master) and the Secondary (Slave) peers should be located. The actual roles performed by the two peers at any given moment may differ from their designations, as a result of a manual role change or an automatic failover.

This command is issued on the Primary peer and affects both peers. For the command to be successfully completed, the HyperSwap relationship has to be activated.

Specifying the new designations is not mandatory. If they are not specified, the command swaps the designations of both peers: the Primary changes to Secondary, and the Secondary changes to Primary.

Example:

```
ha_change_designation vol=regular_volume new_designation=Secondary
```

Output:

```
Command 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_IS_NOT_HA**

  This is not a high availability volume.
- **CONS_GROUP_BAD_NAME**

  Consistency Group name does not exist.
- **CONS_GROUP_IS_NOT_HA**

Local Consistency Group does not have HA 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.
- **HA_IS_NON_OPERATIONAL**

  Ha is non-operational
- **VOLUME_BELONGS_TO_HA_CONS_GROUP**

  This volume belongs to a Ha consistency group. The operation can be carried out on a volume which is not part of a consistency group, or on the consistency group itself.
- **TARGET_BAD_NAME**

  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
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  Data reduced tier is offline, operation is not allowed

  Troubleshooting: Contact IBM Support

Enabling automatic failover in a HyperSwap relationship

Use the **ha_high_availability_enable** command to enable automatic failover in a HyperSwap relationship.

```
ha_high_availability_enable < vol=VolName | cg=cgName >
```

Parameters

Name	Type	Description	Mandatory
vol	Object name	Master volume.	N
cg	Object name	Master consistency group name or a list of master consistency groups.	N

Automatic failover must be enabled if it was disabled in the past. The actual state of the automatic failover, as reported in **ha_list**, depends on several factors, including the user enablement and data synchronization state. It is possible that even though automatic failover is enabled by this command, it remains inactive due to other factors.

The command requires that the HyperSwap relationship be active. It updates both peers, but must be issued only on the Master peer volume or consistency group.

Example:

```
ha_high_availability_enable vol=regular_volume
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- **VOLUME_BAD_NAME**

  Volume name does not exist
- **CONS_GROUP_BAD_NAME**

  Consistency Group name does not exist.
- **VOLUME_IS_NOT_HA**

  This is not a high availability volume.
- **CONS_GROUP_IS_NOT_HA**

  Local Consistency Group does not have HA definitions.
- **TARGET_NOT_CONNECTED**

  There is currently no connection to the target system
- **REMOTE_TARGET_NOT_CONNECTED**

  There is currently no connection from the target system
- **LOCAL_PEER_IS_NOT_MASTER**

  Local peer is not the master
- **VOLUME_BELONGS_TO_HA_CONS_GROUP**

  This volume belongs to a Ha consistency group. The operation can be carried out on a volume which is not part of a consistency group, or on the consistency group itself.
- **HA_IS_NOT_OPERATIONAL**

  This HA relation is not operational. The operation cannot be carried out on a non-operational HA relation.
- **HA_HIGH_AVAILABILITY_IS_ALREADY_ENABLED**

  The Ha high availability ability is already enabled (valid only for ha_high_availability_enable).
- **HA_PEER_QUORUM_WITNESS_CONFIGURATION_NOT_VERIFIED**

  The peer QW connectivity is not verified
- **HA_REMOTE_PEER_QUORUM_WITNESS_CONFIGURATION_NOT_VERIFIED**

  The remote peer QW connectivity is not verified

Disabling automatic failover in a HyperSwap relationship

Use the **ha_high_availability_disable** command to disable automatic failover in a HyperSwap relationship.

```
ha_high_availability_disable < vol=VolName | cg=cgName >
```

Parameters

Name	Type	Description	Mandatory
vol	Object name	Master volume.	N
cg	Object name	Master consistency group name or a list of master consistency groups.	N

The user may choose to disable automatic failover and thereby prevent it in some maintenance scenarios, notably when the Quorum Witness is being replaced.

The command updates both peers, but must be issued only on the Master peer volume or consistency group.

Example:

```
ha_high_availability_disable vol=regular_volume
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- **VOLUME_BAD_NAME**

  Volume name does not exist
- **CONS_GROUP_BAD_NAME**

  Consistency Group name does not exist.
- **VOLUME_IS_NOT_HA**

  This is not a high availability volume.
- **CONS_GROUP_IS_NOT_HA**

  Local Consistency Group does not have HA definitions.
- **TARGET_NOT_CONNECTED**

There is currently no connection to the target system

- **REMOTE_TARGET_NOT_CONNECTED**

  There is currently no connection from the target system
- **LOCAL_PEER_IS_NOT_MASTER**

  Local peer is not the master
- **VOLUME_BELONGS_TO_HA_CONS_GROUP**

  This volume belongs to a Ha consistency group. The operation can be carried out on a volume which is not part of a consistency group, or on the consistency group itself.
- **HA_IS_NOT_OPERATIONAL**

  This HA relation is not operational. The operation cannot be carried out on a non-operational HA relation.
- **HA_HIGH_AVAILABILITY_IS_ALREADY_DISABLED**

  The Ha high availability ability is already disabled (valid only for ha_high_availability_disable).

Converting a HyperSwap relationship into a sync mirror

Use the **ha_convert_into_mirror** command to change a HyperSwap relationship into a sync mirror.

```
ha_convert_into_mirror < vol=VolName | cg=cgName >
```

Parameters

Name	Type	Description	Mandatory
vol	Object name	Master volume.	N
cg	Object name	Master consistency group name or a list of master consistency groups.	N

Pre-requisites:

- The Slave volume is not mapped
- The local peer role is Master
- If carried out on a Slave volume, the HyperSwap connectivity must be down (verified with the **target_list** command)

The operation modifies the Slave volume SCSI identity. The change affects the WWN and the serial, while the volume name, external ID, and metadata remain unchanged.

Once the HyperSwap relationship is converted into a sync mirror, it will no longer be retrieved by the **ha_list** command. Instead, it will appear in the output of the **mirror_list** command.

Example:

```
ha_convert_into_mirror vol=regular_volume
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- **VOLUME_BAD_NAME**

  Volume name does not exist
- **CONS_GROUP_BAD_NAME**

  Consistency Group name does not exist.
- **VOLUME_IS_NOT_HA**

  This is not a high availability volume.
- **CONS_GROUP_IS_NOT_HA**

  Local Consistency Group does not have HA definitions.
- **TARGET_NOT_CONNECTED**

  There is currently no connection to the target system
- **REMOTE_TARGET_NOT_CONNECTED**

  There is currently no connection from the target system
- **LOCAL_PEER_IS_NOT_MASTER**

  Local peer is not the master
- **SLAVE_VOLUME_IS_MAPPED**

  The Slave volume is mapped
- **VOLUME_BELONGS_TO_HA_CONS_GROUP**

  This volume belongs to a Ha consistency group. The operation can be carried out on a volume which is not part of a consistency group, or on the consistency group itself.
- **HA_IS_NOT_OPERATIONAL**

  This HA relation is not operational. The operation cannot be carried out on a non-operational HA relation.

Converting a sync mirror into a HyperSwap relationship

Use the **mirror_convert_into_ha** command to change a sync mirror into a HyperSwap relationship.

```
mirror_convert_into_ha < vol=VolName | cg=cgName >
```

Parameters

Name	Type	Description	Mandatory
vol	Object name	Master volume.	N
cg	Object name	Master consistency group name or a list of master consistency groups.	N

Prerequisites:

- The remote volume is not mapped
- The mirror is a sync mirror
- The local mirror is the Master
- The target connectivity is active and sufficient for HyperSwap (verified with the **target_list** command)
- The peers are connected to a properly configured and active Quorum Witness
- Both volumes have the same name

The command is issued on the Master peer and affects both peers, that are converted to HyperSwap at the same time. The HyperSwap functionality requires certain configuration elements, such as Quorum Witness, and will fail if they do not exist or are not in the required state.

Once the command has successfully completed, the relationship will no longer be retrieved by the **mirror_list** command, but rather by the **ha_list** command.

The conversion process usually continues after the command completion event has been issued. While it continues, the indication of the high availability activation state *Automatic Failover* is *Inactive*, because metadata is still not synchronized. Once the conversion has completed, the HyperSwap volume can be mapped from the remote system and the new mapping can be used to perform I/O to that volume.

Example:

```
mirror_convert_into_ha vol=regular_volume
```

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
- **BAD_REMOTE_VOLUME_NAME**

  Slave 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
- **MIRROR_IS_STANDBY**
  mirror is marked as standby
- **TARGET_NOT_CONNECTED**
  There is currently no connection to the target system
- **REMOTE_TARGET_NOT_CONNECTED**
  There is currently no connection from the target system
- **LOCAL_PEER_IS_NOT_MASTER**
  Local peer is not the master
- **SLAVE_VOLUME_IS_MAPPED**
  The Slave volume is mapped
- **VOLUME_BELONGS_TO_MIRRORED_CONS_GROUP**
  Volume mirror is part of Consistency Group mirror.
- **TARGET_HAS_NO_QUORUM_WITNESS**
  The local target does not have a quorum witness defined.
- **REMOTE_TARGET_HAS_NO_QUORUM_WITNESS**
  The target on the remote system does not have a quorum witness defined.
- **HA_LOCAL_PEER_HAS_NO_QUORUM_WITNESS_CONNECTIVITY**
  The local peer connectivity to the QW is not operational.
- **HA_INCOMPATIBLE_TARGET_VERSION**
  HA is not supported between the system versions of the specified peers.
- **MIRROR_TYPE_IS_NOT_SYNC**
  Mirror type is not Sync Best Effort.
- **MIRROR_IS_NON_OPERATIONAL**
  Mirror is non-operational
- **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_MAX_HA_REACHED**
  Maximum number of HA already defined on remote machine
- **LOCAL_MAX_HA_REACHED**
  Maximum number of HA already defined on local machine
- **VOLUME_HAS_OLVM**
  IBM Hyper-Scale Mobility relation is defined for this volume
- **VOLUME_HAS_HA**
  This operation is forbidden on a volume with HA relation.
- **MAX_NUM_OF_PROXY_VOLUME_REACHED**
  Failed to create mirror, max number of proxy volumes exceeded.
- **TARGET_VOLUME_HAS_OLVM**

This target volume is part of an IBM Hyper-Scale Mobility relation.

- **TARGET_VOLUME_HAS_HA**
  This operation is forbidden with a target volume with HA relation.
- **REMOTE_MAX_METADATA_OBJECTS_REACHED**
  Maximal number of metadata objects has been reached on a remote system.
- **LOCAL_AND_REMOTE_VOLUME_NAMES_ARE_DIFFERENT**
  Local and remote volume names are different.
- **HA_TARGET_QUORUM_WITNESS_IS_NOT_ACTIVATED**
  The quorum witness associated with the target is deactivated.
- **TARGET_PEER_NOT_HEALTHY**
  Target peer health not marked as healthy
  Troubleshooting: check quorum witness configuration
- **HA_CONNECTIVITY_NOT_SUFFICIENT**
  The connectivity between the systems is not sufficient for HA.
- **HA_PEER_QUORUM_WITNESS_CONFIGURATION_NOT_VERIFIED**
  The peer QW connectivity is not verified
- **HA_REMOTE_PEER_QUORUM_WITNESS_CONFIGURATION_NOT_VERIFIED**
  The remote peer QW connectivity is not verified
- **HOST_TYPE_IS_NOT_CONFIGURED**
  Cannot associate a hyperswap volume with a host with unconfigured type. IMPORTANT: Please read the Hyperswap chapter in the 'Best Practice' document to understand the solution requirements.

Creating a new Quorum Witness

Use the **quorum_witness_define** command to create a new Quorum Witness definition in the system and connect the system to the Quorum Witness.

```
quorum_witness_define name=qw_name certificate=qw_certificate address=qw_address
[ port=qw_port ] [ activate=<yes|no> ]
```

Parameters

Name	Type	Description	Mandatory	Default
name	Object name	The name of the Quorum Witness to be created.	Y	N/A
certificate	N/A	The public certificate of the Quorum Witness.	Y	N/A
address	N/A	The Quorum Witness address: IPv4, IPv6 (full format only) or DNS name.	Y	N/A
port	Positive integer	The port used for Quorum Witness communications.	N	8460
activate	Boolean	Defines whether to activate the Quorum Witness upon creation.	N	yes

This command defines a Quorum Witness to be used for HyperSwap relationships. Up to two Quorum Witness instances may be defined. Setting the **activate** parameter to *yes* will also activate the Quorum Witness and start the communication between the system and the Quorum Witness.

Example:

```
quorum_witness_define name=q1 certificate="Certificate" address=195.7.15.2
```

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

- **QUORUM_WITNESS_NAME_ALREADY_EXISTS**

  Quorum witness name already exists.
- **QUORUM_WITNESS_ADDRESS_ALREADY_EXISTS**

  The quorum witness IPv4/6 address or DNS name with same port number already exists.

  Troubleshooting: Use the already configured quorum witness or change the address or port number.
- **QUORUM_WITNESS_MISSING_SERVICE_CERTIFICATE**

  No valid certificate is defined for the quorum witness service.

  Troubleshooting: use PKI commands to define certificate for the quorum witness service
- **MAX_QUORUM_WITNESSES_REACHED**

  Too many quorum witness instances are defined, delete one or more quorum witness instances before retrying.
- **MAX_ACTIVE_QUORUM_WITNESSES_REACHED**

  The maximum number of active quorum witness instances was reached.

  Troubleshooting: Deactivate an active quorum witness and then try again.
- **SSL_CERTIFICATE_HAS_EXPIRED**

  SSL certificate has expired.
- **SSL_CERTIFICATE_ISSUER_NOT_FOUND**

  SSL certificate issuer not found in certificate chain.
- **SSL_CERTIFICATE_INVALID_FORMAT**

  SSL certificate format is invalid or corrupted.
- **SSL_CERTIFICATE_CHAIN_EMPTY**

No certificates found in input.

- **SSL_CERTIFICATE_VERIFICATION_INTERNAL_ERROR**
  SSL certificate verification has failed because of internal system error.
- **SSL_CERTIFICATE_VERIFICATION_FAILED**
  SSL certificate chain verification failed.
- **SSL_CERTIFICATE_NOT_YET_VALID**
  SSL certificate is not yet valid.

Listing Quorum Witnesses

Use the **quorum_witness_list** command to list all Quorum Witnesses defined in the system, or only the specified one.

```
quorum_witness_list [ name=qw_name ] [ domain=DomainName ]
```

Parameters

Name	Type	Description	Mandatory	Default
name	Object name	The internal name of the Quorum Witness.	N	All Quorum Witnesses
domain	Object name	The domain name.	N	All domains

The command output includes the following fields:

- **name**: The Quorum Witness internal name
- **quorum_id**: A globally unique Quorum Witness ID
- **address**: The v6 or DNS address for communicating with the Quorum Witness
- **port**: The port for communicating with the Quorum Witness
- **state**: The state of the Quorum Witness in the system. The following values are available: `Activating`, `Activated`, `Deactivating`, and `Deactivated`.
- **connection**: The state of the connection with the Quorum Witness, accumulated across all Quorum Node connection statuses.
- **external_name**: The Quorum Witness external name
- **db_health**: The state of the Quorum Witness DB health. The following values are available:

Value	Meaning
`OK`	
`Recovery needed`	A problem was identified in the Quorum Witness DB, that prevents it from operating properly.
`Restore pending`	The Quorum Witness DB was initialized. Once the Quorum Witness is activated, the system will automatically re-register to the Quorum Witness, and restore the information pertaining to its HyperSwap relationships.

The following optional fields can be listed by explicitly specifying the proper columns:

- **heartbeats_ok**: Indicates whether heartbeats between the system and the Quorum Witness are properly sent and received

- **secure_connection**: The state of the security established on the connection. The possible values are:

Value	Meaning
None	Security is disabled on the Quorum Witness connectivity.
Verified	The connection security is verified.
Unauthorized	The system certificate was rejected by the Quorum Witness.

- **version**: The Quorum Witness software version
- **id**
- **certificate**
- **db_init**: The timestamp (in mono-time) of the last DB initialization
- **first_event_id**: The ID of the first relevant event
- **last_event_id**: The ID of the last relevant event

Example:

```
quorum_witness_list name=q1
```

Output:

```
Name     ID                 Address     Port   State       Connection    External Name
------   ----------------   ----------  -----  ----------  ----------    -------------
q1       quorum_witness_id  195.7.15.2  4567   Activated   Up            quorum_witness

Cont.:

Db Health
---------
OK
```

Field ID	Field output	Default position
name	Name	1
quorum_id	ID	2
address	Address	3
port	Port	4
state	State	5
connection	Connection	6
external_name	External Name	7
db_health	Db Health	8
heartbeats_ok	Heartbeating	N/A
use_secure	Use Secure	N/A
secure_connection	Secure Connection	N/A
version	Version	N/A
id	ID	N/A
certificate	Certificate	N/A
db_init	DB Init	N/A
first_event_id	First Event Id	N/A
last_event_id	Last Event Id	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Updating a Quorum Witness definition

Use the **quorum_witness_update** command to update a Quorum Witness definition.

```
quorum_witness_update name=qw_name [ certificate=new_qw_certificate ]
 [ address=new_qw_address ] [ port=new_qw_port ]
```

Parameters

Name	Type	Description	Mandatory	Default
name	Object name	The Quorum Witness internal name.	Y	N/A
certificate	N/A	A new public certificate of the Quorum Witness.	N	Current value.
address	N/A	The Quorum Witness address: IPv4, IPv6 (full format only) or DNS name.	Y	N/A
port	Positive integer	A new communication port of the Quorum Witness.	N	Current value.

As a prerequisite for successfully completing this command, the Quorum Witness must be deactivated.

Example:

```
quorum_witness_update name=q1 address=192.6.10.7
```

Output:

```
Command executed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed

User Category	Permission
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- **QUORUM_WITNESS_BAD_NAME**

  Quorum witness name does not exist.
- **QUORUM_WITNESS_ADDRESS_ALREADY_EXISTS**

  The quorum witness IPv4/6 address or DNS name with same port number already exists.

  Troubleshooting: Use the already configured quorum witness or change the address or port number.
- **CANT_UPDATE_ACTIVATED_QUORUM_WITNESS**

  Cannot update an activated quorum witness.
- **SSL_CERTIFICATE_HAS_EXPIRED**

  SSL certificate has expired.
- **SSL_CERTIFICATE_ISSUER_NOT_FOUND**

  SSL certificate issuer not found in certificate chain.
- **SSL_CERTIFICATE_INVALID_FORMAT**

  SSL certificate format is invalid or corrupted.
- **SSL_CERTIFICATE_CHAIN_EMPTY**

  No certificates found in input.
- **SSL_CERTIFICATE_VERIFICATION_INTERNAL_ERROR**

  SSL certificate verification has failed because of internal system error.
- **SSL_CERTIFICATE_VERIFICATION_FAILED**

  SSL certificate chain verification failed.
- **SSL_CERTIFICATE_NOT_YET_VALID**

  SSL certificate is not yet valid.

Renaming a Quorum Witness

Use the **quorum_witness_rename** command to rename a Quorum Witness.

```
quorum_witness_rename name=qw_name new_name=qw_new_name
```

Parameters

Name	Type	Description	Mandatory
name	Object name	The Quorum Witness internal name.	Y
new_name	Object name	A new internal name of the Quorum Witness.	Y

Example:

```
quorum_witness_rename name=q1 new_name=q2
```

Output:

```
Command 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

- **QUORUM_WITNESS_BAD_NAME**
  Quorum witness name does not exist.
- **QUORUM_WITNESS_NAME_ALREADY_EXISTS**
  Quorum witness name already exists.

Deleting a Quorum Witness

Use the **quorum_witness_delete** command to delete a Quorum Witness.

```
quorum_witness_delete name=qw_name
```

Parameters

Name	Type	Description	Mandatory
name	Object name	The Quorum Witness internal name.	Y

The command can be issued only for a Quorum Witness that is not in use (attached to any target) and is deactivated or is being deactivated, that is its state is *Deactivated* or *Deactivating*.

Example:

```
quorum_witness_delete name=q1 -y
```

Output:

```
Command executed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed

User Category	Permission
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_DELETE_QUORUM_WITNESS**
  Are you sure you want to delete the quorum witness *Quorum Witness*?

Return codes

- **QUORUM_WITNESS_BAD_NAME**
  Quorum witness name does not exist.
- **CANT_DELETE_AN_ACTIVATED_QUORUM_WITNESS**
  Cannot delete an activated quorum witness.
  Troubleshooting: Deactivate the quorum witness and then try again.
- **CANT_DELETE_QUORUM_WITNESS_IN_USE**
  Cannot delete a quorum witness when it is in use by a target.

Activating a Quorum Witness

Use the **quorum_witness_activate** command to activate a Quorum Witness.

```
quorum_witness_activate name=qw_name
```

Parameters

Name	Type	Description	Mandatory
name	Object name	The Quorum Witness internal name.	Y

This commands activates a Quorum Witness and starts heartbeat and status communication between the system and the Quorum Witness. The command is asynchronous, its completion is tracked by return codes.

Example:

```
quorum_witness_activate name=q1
```

Output:

```
Command executed 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

- **MAX_ACTIVE_QUORUM_WITNESSES_REACHED**

  The maximum number of active quorum witness instances was reached.

  Troubleshooting: Deactivate an active quorum witness and then try again.
- **QUORUM_WITNESS_BAD_NAME**

  Quorum witness name does not exist.
- **QUORUM_WITNESS_IS_ALREADY_ACTIVE**

  The quorum witness is already active or is being activated.

  Troubleshooting: Wait for the activation process to complete. If this issue persists, contact IBM support.
- **QUORUM_WITNESS_MISSING_SERVICE_CERTIFICATE**

  No valid certificate is defined for the quorum witness service.

  Troubleshooting: use PKI commands to define certificate for the quorum witness service
- **SSL_CERTIFICATE_HAS_EXPIRED**

  SSL certificate has expired.

Deactivating a Quorum Witness

Use the **quorum_witness_deactivate** command to deactivate a Quorum Witness.

```
quorum_witness_deactivate name=qw_name
```

Parameters

Name	Type	Description	Mandatory
name	Object name	The Quorum Witness internal name.	Y

This commands deactivates a Quorum Witness and stops the heartbeat and status communication between the system and the Quorum Witness. It is not allowed to deactivate a Quorum Witness that is in use by an active HyperSwap relationship.

Example:

```
quorum_witness_deactivate name=q1 -y
```

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

Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_DEACTIVATE_QUORUM_WITNESS**

  Are you sure you want to deactivate quorum witness *Quorum Witness*?

Return codes

- **QUORUM_WITNESS_BAD_NAME**

  Quorum witness name does not exist.
- **QUORUM_WITNESS_IS_ALREADY_INACTIVE**

  The quorum witness is already inactive or is being deactivated.

  Troubleshooting: Wait for the deactivation process to complete. If this issue persists, contact IBM support.
- **CANT_DEACTIVATE_QUORUM_WITNESS_IN_USE**

  Can't deactivate a quorum witness that is in use by a target with an active high availability relation.

Listing the Quorum Witness connection status

Use the **quorum_witness_connections_list** command to list the connection status of each module with the Quorum Witness or only of the specified one.

```
quorum_witness_connections_list [ name=qw_name ] [ module=module_id ]
```

Parameters

Name	Type	Description	Mandatory	Default
name	Object name	The Quorum Witness internal name.	N	All quorums connections.
module	N/A	The module ID.	N	All modules connections.

The command output includes the following fields:

- **module_id**: The ID of the module on which the Quorum Witness is running
- **name**: The Quorum Witness internal name
- **connection**: The state of the connection with the Quorum Witness. The possible values are:

Value	Meaning
Connected	A Quorum Node is successfully connected to the Quorum Witness.
Failed to initialize	HTTP connection initialization failed.
Malformed URL	

Value	Meaning
Cannot resolve the proxy address	
Cannot resolve the host address	
Cannot connect to peer	
HTTP communication error	
Out of memory	
Connection timeout	The Quorum Witness may be down or the address/port pair is incorrect.
HTTP Post communication error	
Secure connection issue	See details in the **secure_connection** field below.
Too many redirects	
Bad content encoding	
Transient error	
Failed to chunk data	
Quorum node has failed	

- **secure_connection**: The state of the security established on the connection. The possible values are:

Value	Meaning
None	Security is disabled on the Quorum Witness connectivity.
Verified	The connection security is verified.
Unauthorized	The system certificate was rejected by the Quorum Witness.
General SSL/TLS failure	
Peer verification failed	
Problem with local certificate	
Secure cipher error	
CA certificate is unusable	The provided Quorum Witness certificate is incorrect.
Failed initializing secure communication	
Could not load CACERT file, missing or wrong format	
Failed to shut down the SSL/TLS connection	
Could not load CRL file, missing or wrong format	
Issuer check failed	

- **heartbeating**: Indicates whether the module heartbeats are successfully processed by the Quorum Witness.

Example:

```
quorum_connection_list name=q1 module=1:Module:3
```

Output:

```
Module ID    Name    Connection   Secure Connection   Heartbeating
------------ ------- ------------ ------------------- --------------
1:Module:3   q1      Connected    Verified            yes
```

Field ID	Field output	Default position
module_id	Module ID	1
name	Name	2
connection	Connection	3
secure_connection	Secure Connection	4
heartbeating	Heartbeating	5

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Allowed

Getting Quorum Witness information

Use the **quorum_witness_info_get** command to get Quorum Witness information from a Quorum Witness.

```
quorum_witness_info_get name=qw_name
```

Parameters

Name	Type	Description	Mandatory
name	Object name	The Quorum Witness internal name.	Y

The command output includes the following fields:

- **internal_name**: The Quorum Witness internal name
- **quorum_id**: A globally unique Quorum Witness ID
- **name**: The Quorum Witness external name
- **version**: The Quorum Witness software version

The following optional fields can be listed by explicitly specifying the proper columns:

- **db_state**: JSON string (up to 256 bytes), representing the DB disk usage
- **host_info**: The type and the version of the host's operating system, on which the Quorum Witness is running
- **protocol**: The protocol version
- **network_load**: JSON string (up to 128 bytes), representing the count of connections and requests
- **cpu_load**: JSON string (up to 128 bytes), representing the CPU consumption percentage
- **last_event_id**: The ID of the last relevant event
- **db_init**: The timestamp (in mono-time) of the last DB initialization

Example:

```
quorum_witness_info_get name=q1
```

Output:

```
Name   ID                                 External Name   Version
------ ---------------------------------- --------------- ---------
q1     7f2f35834fea48b6b335aa09fd9179ae   FVT-QW104       1.0.0
```

Field ID	Field output	Default position
internal_name	Name	1
quorum_id	ID	2
name	External Name	3
version	Version	4
db_state	DB state	N/A
host_info	Host	N/A
protocol	Protocol Version	N/A
network_load	Network Load	N/A
cpu_load	CPU Load	N/A
last_event_id	Last Event ID	N/A
db_init	DB Initialization Timestamp	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Return codes

- **QUORUM_WITNESS_BAD_NAME**

  Quorum witness name does not exist.

- **QUORUM_WITNESS_RESPONSE_TIMEOUT**

  No response from quorum witness *Quorum Witness* within given timeout window.

Chapter 12. 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.

```
dm_activate vol=VolName
```

Parameters

Name	Type	Description	Mandatory
vol	Object name	The destination volume for data migration activation.	Y

This command activates the data migration process. This is either an initial activation or an activation after deactivation.

Upon activation, the data migration is tested in the same way as when using **dm_test** (see Testing the data migration definition), and this command fails if the data migration test fails.

This command has no effect if the process is already active.

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- **VOLUME_BAD_NAME**
  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

© Copyright IBM Corp. 2017

- **REMOTE_VOLUME_NO_WRITE_ACCESS**
  Remote volume is write protected
- **BAD_REMOTE_VOLUME_SIZE**
  Master and slave volumes contain a different number of blocks
- **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

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
- **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

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

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_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_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

- **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
- **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.
- **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
- **VOLUME_TOO_BIG**
  No space to allocate volume

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	
force_delete	Boolean	When set to *yes*, forces the deletion even if data migration is not complete.	N	no

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
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
- **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
- **DM_OPERATION_NOT_ALLOWED_NEITHER_SOURCE_NOR_TARGET_CONTAIN_FULL_DATA**

  Data Migration was setup without source update and synchronization did not complete yet. As a result, neither the source volume nor the target volume contain a full image of the data.

Listing data migration statuses

Use the **dm_list** command to list data migration configuration and status.

```
dm_list [ vol=VolName ] [ domain=DomainName ]
```

Parameters

Name	Type	Description	Mandatory	Default
vol	Object name	The name of the volume to be listed.	N	All data migration volumes.
domain	Object name	The domain name.	N	All Domains

This command lists all data migration configuration and statuses, including the following information:

- Volume name
- Target name
- LUN
- Volume size (GB)
- Migration completed (GB)
- Migration activation (active/inactive)
- Migration status (synchronized, unsynchronized)
- Migration remaining (GB)
- Migration remaining (%)
- Estimated time to completion

Field ID	Field output	Default position
local_volume_name	Local Volume	1
target_name	Remote System	2
remote_volume_lun	Remote LUN	3
active	Active	4
sync_state	Status	5
connected	Target Connected	6
size_to_synchronize	Size To Sync (MiB)	N/A
operational	Operational	N/A
sync_progress	Sync Progress (%)	N/A
start_migration_automatically	Start Data Migration Automatically	N/A

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
- **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

Chapter 13. 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.

```
olvm_create < vol=VolName remote_pool=RemotePoolName > target=TargetName
```

Parameters

Name	Type	Description	Mandatory
vol	Object name	An IBM Hyper-Scale Mobility volume on the local system.	Y
target	Object name	Remote system containing the destination volume.	Y
remote_pool	Object name	Name of the storage pool to contain the destination volume.	Y

This command creates an IBM Hyper-Scale Mobility relation through identifying the source volume and the destination system and storage pool.

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- **VOLUME_SIZE_VERY_LARGE_ARE_YOU_SURE**

  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**

© Copyright IBM Corp. 2017

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.
- **VOLUME_TOO_BIG**
  No space to allocate volume
- **VOLUME_SIZE_ABOVE_LIMIT**
  Volume size specified is above limit
- **INVALID_SLICE_OFFSET**
  Slice offset is illegal
- **VOLUME_IS_OLVM_PROXY**
  The volume is in an IBM Hyper-Scale Mobility Proxy phase.
- **VOLUME_IS_SLAVE**
  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
- **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.
- **REMOTE_VOLUME_HAS_MIRRORING_SNAPSHOTS**
  Remote Volume has snapshots created by previous mirroring process.
- **VOLUME_HAS_MULTIPLE_MIRRORS**
  volume has multiple mirrors, operation not allowed or target must be specified

- **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
- **SYSTEM_OUT_OF_PHYSICAL_SPACE**

  Operation not allowed while the system is out of physical space.

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
- **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
- **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

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
- **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

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
- **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

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.
- **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
- **MAX_METADATA_OBJECTS_REACHED**

  Maximal number of metadata objects has been reached.
- **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

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_HAS_MULTIPLE_MIRRORS**
  volume has multiple mirrors, operation not allowed or target must be specified
- **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

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

Chapter 14. 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.

```
custom_event description=Description
[ severity=<INFORMATIONAL|WARNING|MINOR|MAJOR|CRITICAL> ]
```

Parameters

Name	Type	Description	Mandatory	Default
description	String	Description of the event.	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:

```
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> ]
```

© Copyright IBM Corp. 2017

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_detailes
```

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

Defining a new event notification destination

Use the **dest_define** command to define a new destination for event notifications.

```
dest_define
dest=DestName type=<SNMP|EMAIL|SMS|HTTPS>
< snmp_manager=SNMPManager | < uri=HTTPSaddress
[ proxy=ProxyAddress [ proxy_port=ProxyPortNum ] ]
> | email_address=email |
<area_code=AreaCode number=PhoneNumber> | user=UserName>
[ smtpgws=<SMTPGW1[,SMTPGW2]...|ALL> | smsgws=<SMSGW1[,SMSGW2]...|ALL> ]
[ heartbeat_test_hour=HH:MM
[ heartbeat_test_days=Day ] ] [ domain=DomainList ]
```

Parameters

Name	Type	Description	Mandatory	Default
dest	Object name	Destination name.	Y	N/A
type	Enumeration	Destination type for event notifications: be email, SMS, HTTPS or SNMP.	Y	N/A
snmp_manager	N/A	IP address or DNS name of the SNMP manager.	N	N/A
uri	N/A	IP address or DNS name of the HTTPS server. If a port different from the default should be used, specify it here.	N	N/A
proxy	N/A	IP address or DNS name of the proxy server to send HTTPS over.	N	None
proxy_port	Integer	Proxy port number to send HTTPS through. The default is 1080.	N	None
email_address	N/A	Email address.	N	N/A
smtpgws	Object name	List of SMTP gateways to be used.	N	ALL (all gateways).
area_code	N/A	Area code of the cellular number for SMS notification. Use digits, '-' or '.'	N	N/A
number	N/A	Cellular number for SMS notification. Use digits, '-' or '.'	N	N/A
smsgws	Object name	SMS gateways to be used for this destination.	N	ALL (all gateways).
user	Object name	User name, where the user's email or phone are used.	N	N/A
heartbeat_test_hour	N/A	The hour for periodic heartbeat testing in the format HH:MM	N	No heartbeat
heartbeat_test_days	N/A	List of days for heartbeat testing: a comma-separated list of 3-letter day names (such as "mon", "mon,fri", etc.).	N	No heartbeat

Name	Type	Description	Mandatory	Default
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
- **SMSGWS_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 by 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 XIV attribute.	N	no
domain	Object name	The domain name.	N	All Domains

This command lists the configuration of all defined destinations, or of a specific destination.

Field ID	Field output	Default position
name	Name	1
type	Type	2
email_address	Email Address	3
area_code	Area Code	4
number	Phone Number	5
snmp_manager	SNMP Manager	6
uri	HTTPS Address	7
gateways	Gateways	N/A
user	User	8
heartbeat_test_days	Heartbeat Days	N/A
heartbeat_test_hour	Heartbeat Time	N/A
creator	Creator	N/A
proxy	proxy server address	N/A
proxy_port	proxy port number	N/A

Example:

```
dest_list
```

Output:

```
Name              Type  Email Address                 Phone Number  Gateways
storagemanager    EMAIL storageadmin@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

Name	Type	Description	Mandatory
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 ]
[ smsgw=SMSGatewayName ] [ internal=<yes|no> ]
```

Parameters

Name	Type	Description	Mandatory	Default
dest	Object name	Name of the destination to be tested.	Y	N/A

Name	Type	Description	Mandatory	Default
management_ip	N/A	Management IP used for sending the event notification.	Y	N/A
smtpgw	Object name	SMTP gateway to be tested.	N	Default system choice.
smsgw	Object name	SMS gateway to be tested.	N	Default system choice.
internal	Boolean	Must be specified for XIV-internal destinations.	N	no

This command tests a destination by sending a test message, SMS or SNMP trap. Note that a successful return code from this command does not ensure notification delivery.

Some problems with SNMP, email, and SMS delivery may fail to be detected.

For email messages, the SMTP gateway must be specified (the destination is only tested through that gateway). The same applies to the SMS gateway.

Access control

User Category	Permission	Condition
Storage administrator	Conditionally Allowed	Allowed, unless the internal parameter is specified.
Storage integration administrator	Disallowed	N/A
Application administrator	Conditionally Allowed	Allowed, unless the internal parameter is specified.
Security administrator	Disallowed	N/A
Read-only users	Disallowed	N/A
Technicians	Allowed	N/A

Return codes

- **DEST_NAME_DOES_NOT_EXIST**
  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
- **SMSGWS_MUST_BE_SPECIFIED_FOR_DEST_TYPE**
  Destination must have SMS gateways
- **SMSGWS_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 ]
[ smsgws=<SMSGW1[,SMSGW2]...|ALL> ]
[ user=UserName ] [ heartbeat_test_hour=HH:MM ]
[ heartbeat_test_days=Day ] [ domain=DomainList ]
```

Parameters

Name	Type	Description	Mandatory	Default
dest	Object name	Destination name.	Y	N/A
snmp_manager	N/A	IP address or DNS name of the SNMP manager.	N	Keep unchanged.
uri	N/A	IP address or DNS name of HTTPS server.	N	Keep unchanged.
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.

Name	Type	Description	Mandatory	Default
area_code	N/A	Area code of the cellular number for SMS notification.	N	Keep unchanged.
number	N/A	Cellular number for SMS notification.	N	Keep unchanged.
smsgws	Object name	SMS gateways to be used.	N	Keep unchanged.
user	Object name	User name, where the user's email or phone are used.	N	Keep unchanged.
heartbeat_test_hour	N/A	The hour of periodic heartbeat testing	N	Keep unchanged.
heartbeat_test_days	N/A	List of days for heartbeat testing: a comma-separated list of 3-letter day names (such as "mon", "mon,fri", and so on).	N	Keep unchanged.

The parameters of this command are identical to the Defining a new event notification destination command, except that the destination type cannot be changed. All relevant fields must be specified (not only the ones that are being changed).

Example:

```
dest_update dest=storagemanager
email_address=admin@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 by included in the destination domains.
- **DESTINATION_IS_NOT_IN_RULE_DOMAINS**
  Destination must by included in the rule domains.
- **DESTINATION_IS_NOT_IN_DESTGROUP_DOMAINS**
  Destinations must by 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 by 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

Name	Type	Description	Mandatory
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 by included in the destination group domains.
- **DOMAIN_DOESNT_EXIST**
  Domain does not exist.
- **DESTGROUP_IS_NOT_IN_RULE_DOMAINS**
  Destination groups must by 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

Field ID	Field output	Default position
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
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|smsgw|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

Name	Type	Description	Mandatory	Default
internal	Boolean	Filter XIV internal events.	N	no filter
beg	Integer	Index of the first event to list. If negative, then counts from the end.	N	1
end	Integer	Index of the last event to list (not inclusive). If negative, then counts from the end.	N	last event + 1
count_all	Boolean	If yes, it scans all the events between beginning and end for computing the number of events meeting the criteria.	N	no
domain	Object name	The domain name.	N	All Domains

This command lists system events according to specified criteria, such as minimum severity, event type, and so on. The event list displays the following information for each event: timestamp, severity, code, user and description.

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
                  System has entered ON state.
xiv_development   Storage Pool of size 171GB was created with name
                  'p1_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 'p1_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
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

Field ID	Field output	Default position
default_thresholds.3	MAJOR(def)	10
default_thresholds.4	CRITICAL(def)	11
thresholds.0	INFORMATIONAL	2
thresholds.1	WARNING	3
thresholds.2	MINOR	4
thresholds.3	MAJOR	5
thresholds.4	CRITICAL	6

Example:

```
event_threshold_list
```

Output:

```
Code                                    INFORMATIONAL   WARNING   MINOR
--------------------------------------- --------------- --------- -------
STORAGE_POOL_SNAPSHOT_USAGE_INCREASED   none             80        90
STORAGE_POOL_VOLUME_USAGE_INCREASED     none             80        90

MAJOR   CRITICAL   INFORMATIONAL(def)   WARNING(def)   MINOR(def)   MAJOR(def)
------- ---------- -------------------- -------------- ------------ ------------
95      none       none                 80             90            95
95      none       none                 80             90            95

CRITICAL(def)
---------------
none
none
```

Access control

User Category	Permission	Condition
Storage administrator	Conditionally Allowed	Allowed, unless the internal parameter is specified.
Storage integration administrator	Disallowed	N/A
Application administrator	Conditionally Allowed	Allowed, unless the internal parameter is specified.
Security administrator	Disallowed	N/A
Read-only users	Conditionally Allowed	Allowed, unless the internal parameter is specified.
Technicians	Allowed	N/A

Activating a rule

Use the **rule_activate** command to activate an event notification rule.

```
rule_activate rule=RuleName
```

Parameters

Name	Type	Description	Mandatory
rule	Object name	The name of the rule to be activated.	Y

This command activates the specified rule. An active rule is matched against events and generates notifications. If the rule is already active, this command has no effect.

Access control

User Category	Permission
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
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

Name	Type	Description	Mandatory	Default
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:

```
xcli -u  -c Nextra1 rule_create rule=critical_alerts min_severity=critical destinations=joh
n-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

User Category	Permission
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

Field ID	Field output	Default position
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

User Category	Permission
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.

Name	Type	Description	Mandatory	Default
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 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

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=SMSGW1
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 **smsgw_delete** command to delete an SMS gateway.

```
smsgw_delete smsgw=SMSGatewayName
```

Parameters

Name	Type	Description	Mandatory
smsgw	Object name	SMS gateway to be deleted.	Y

A gateway cannot be deleted if it is part of a notification rule or if it is being used by a destination.

Before deleting an SMS gateway, make sure that all alerting events are cleared.

Example:

```
smsgw_delete smsgw=external-SMSGW
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_DELETE_SMS_GATEWAY**
  Are you sure you want to delete SMS gateway *Gateway*?

Return codes

- **CANNOT_CHANGE_EVENT_CONF_WITH_ALERTING_EVENTS**
  Cannot change 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 **smsgw_list** command to list SMS gateways.

```
smsgw_list [ smsgw=SMSGatewayName ]
```

Parameters

Name	Type	Description	Mandatory	Default
smsgw	Object name	Name of SMS gateway to list.	N	All gateways.

The command lists all SMS gateways, or a specific one. For each SMS gateway, all of its configuration information is listed.

Field ID	Field output	Default position
name	Name	1
email_address	Email Address	2
gateways	SMTP Gateways	3
subject_line	Subject Line	N/A
email_body	Email Body	N/A
priority	Priority	N/A

Example:

```
smsgw_list
```

Output:

```
Name    Email Address                                      SMTP Gateways
SMSGW1  {areacode}{number}@sms2emailserver.yourcompany.com  all
SMSGW2  {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 **smsgw_prioritize** command to set the priorities of the SMS gateways for sending SMS messages.

```
smsgw_prioritize order=<gw1[,gw2]...>
```

Parameters

Name	Type	Description	Mandatory
order	Object name	List of all SMS gateways ordered by priority.	Y

SMS messages can be sent to cell phones through one of the email-to-SMS gateways in this list. This command determines the order in which the storage system attempts to use these SMS gateways.

Only one gateway is used and subsequent gateways are only tried if the preceding ones in this priority list return an error.

Specific SMS destinations may define their own SMS gateways to be used when sending SMS to these destinations, regardless of this list.

Example:

```
smsgw_prioritize order=SMSGW1,SMSGW2
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- **CANNOT_CHANGE_EVENT_CONF_WITH_ALERTING_EVENTS**
  Cannot change 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 **smsgw_rename** command to rename an SMS gateway.

```
smsgw_rename smsgw=SMSGatewayName new_name=Name
```

Parameters

Name	Type	Description	Mandatory
smsgw	Object name	SMS gateway to be renamed.	Y
new_name	Object name	New name for the SMS gateway.	Y

Before renaming an SMS gateway, make sure that all alerting events are cleared.

Example:

```
smsgw_rename smsgw=SMSGW2 new_name=external-SMSGW
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- **CANNOT_CHANGE_EVENT_CONF_WITH_ALERTING_EVENTS**
  Cannot change 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 **smsgw_update** command to update an SMS gateway.

```
smsgw_update smsgw=SMSGatewayName [ email_address=email ]
[ subject_line=SubjectLineScheme ] [ email_body=EmailBodyScheme ]
[ smtpgw=<SMTPGW1[,SMTPGW2]...|ALL> ]
```

Parameters

Name	Type	Description	Mandatory	Default
smsgw	Object name	SMS gateway name.	Y	N/A
email_address	Token String	Format for email address.	N	Leave unchanged.
subject_line	Token String	Format for subject line.	N	Leave unchanged.
email_body	Token String	Format for the email's body.	N	Leave unchanged.
smtpgw	Object name	List of SMTP gateways to be used.	N	The SMTP gateways defined in the smtpgw_prioritize command.

This command updates the configuration information of an existing SMS gateway. For the exact description and documentation of each parameter, see the documentation of Defining an SMS gateway.

This command cannot be executed while there are uncleared alerting events.

Parameters that are not specified will not be changed.

Example:

```
smsgw_update smsgw=SMSGW1
email_address={areacode}{number}@sms2emailserver.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_address	N/A	The reply to address used for outgoing emails sent through this SMTP server.	N	DEFAULT (system-wide reply-to address that applies to all servers).

Several email gateways can be defined to enable notification of events by email or sending SMS messages via email-to-SMS gateways. By default, the system attempts to send each email notification through the first gateway according to the order that you specify. Subsequent gateways are only tried if the first in line returns an error. A specific email destination, or a specific SMS gateway may be defined to use only specific SMTP gateways.

The SMTP protocol dictates that every email message must specify the email address of the sender. This sender address must be a valid address for two reasons:

- Many SMTP gateways require a valid sender address, otherwise they will not forward the email, as a security measure in order to prevent unauthorized usage of the SMTP server. Often this sender address must be limited to a specific domain.
- The sender's address is used as the destination for error messages generated by the SMTP gateways, such as: incorrect email address, full email mailbox and so on.

If the sender's address is not specified for a specific SMTP gateway, a global system-wide sender's address specified in Setting configuration parameters is used.

The user can also configure a reply-to address which is different from the sender's address, if it is required that the return emails be sent to another destination.

Example:

```
smtpgw_define smtpgw=mailserver1 address=smtp.yourcompany.com
from_address=nextra@yourcompany.com
reply_to_address=nextraerrors@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

Field ID	Field output	Default position
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.
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> ]
```

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.

Name	Type	Description	Mandatory	Default
reply_to_address	N/A	The reply-to address used for outgoing emails sent through this SMTP server, or DEFAULT for the system-wide default.	N	Leave unchanged.

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@yurcompany.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
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
- **ONLY_TECHNICIAN_CAN_REFER_TO_INTERNAL_EVENT_OBJECTS**
  Only technician can refer to internal event objects

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

Chapter 15. IP configuration commands

This section describes the command-line interface (CLI) for IP configuration.

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
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.

© Copyright IBM Corp. 2017

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
- **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
- **ILLEGAL_COMPONENT_ID**
  Component ID is illegal
- **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

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
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.

Name	Type	Description	Mandatory	Default
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
passkey	N/A	Secret password.	N	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	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

User Category	Permission
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

Name	Type	Description	Mandatory	Default
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

Name	Type	Description	Mandatory	Default
left_ip	N/A	Lists all IPSec tunnels from this left IP.	N	IPsec tunnels from any left IP
right_ip	N/A	Lists all IPSec tunnels from this right IP.	N	IPsec tunnels to any right IP
module	N/A	Limits the listing to a specific module.	N	All modules

Field ID	Field output	Default position
name	IPSec Connection	1
type	Type	2
status	Status	3
left	Left Interface	4
left_ip	Left Address	5
right_ip	Right Address	6
module	Module	7

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Connecting to a support center

Use the **support_center_connect** command to connect to a support center.

```
support_center_connect
[ < timeout=Timeout [ idle_timeout=IdleTimeout ] > | always_on=<yes|no> ]
[ module=ModuleNumber ] [ password=Password ]
```

Parameters

Name	Type	Description	Mandatory	Default
timeout	N/A	Specifies the duration of the session. After the duration elapses, the session will be disconnected. Time is specified in hh:mm format.	N	none

Name	Type	Description	Mandatory	Default
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
- **REMOTE_SUPPORT_CLIENT_CONNECT_ON_RESTART_IS_RUNNING**
  The Remote Support Client in connect-on-restart mode is running
  Troubleshooting: Run support_center_disconnect to stop it.

Defining a support center

Use the **support_center_define** command to define a support center.

```
support_center_define support_center=SupportCenterName address=Address [ port=port ]
 [ priority=priority ]
```

Parameters

Name	Type	Description	Mandatory	Default
support_center	Object name	The name of the support center server	Y	N/A
address	N/A	The IP address of the support center server	Y	N/A
port	Positive integer	The TCP port to connect to on the support center	N	22
priority	N/A	The priority of the support center (support centers with a higher priority will be connected first)	N	0

Example:

```
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
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

User Category	Permission
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   Connect-on-restart active
------------- --------------------------
              no
```

Field ID	Field output	Default position
state	State	1
connected_support_sessions	Connected sessions	2
minutes_to_timeout	Timeout (min)	3
running_from_module	Module	4
start_time	Connected since	5
destination	Destination	6
automatically_connect_mode	Auto Connect Active	7
stop_automatically_connect	Stop support center automatically connect	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Allowed

Defining a support center connection automatically

Use the **support_center_config** command to configure the connection to a support center automatically.

```
support_center_config automatically_connect=<yes|no> [ connect_through_module1=module ]
 [ connect_through_module2=module ] [ connect_through_module3=module ] [ password=Password ]
```

Parameters

Name	Type	Description	Mandatory	Default
automatically_ connect	Boolean	Enables the automatic connection to the support center.	Y	N/A
connect_through_ module1	N/A	The first module from which to establish a connection to the support center automatically.	N	Module with first management port
connect_through_ module2	N/A	The second module from which to establish a connection to the support center automatically.	N	Module with second management port
connect_through_ module3	N/A	The third module from which to establish a connection to the support center automatically.	N	Module with third management port
password	String	A password set by the customer, that needs to be submitted by support services, in order to start a remote support session. Password format: case-insensitive string of 6-12 alphanumeric characters.	N	none

Example:

```
support_center_config automatically_connect=yes
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Allowed

Return codes

- **MODULE_HAS_NO_SUPPORT_CENTER_PORT**

  The specified module does not have a port from which 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 the automatic connection to a support center

Use the **support_center_config_list** command to display the configuration of the automatic connection to a support center.

```
support_center_config_list
```

Example:

```
support_center_config_list
```

Output:

```
Enable Auto Conn   First Module   Second Module   Third Module
------------------ -------------- --------------- --------------
yes                1              2               -1
```

Field ID	Field output	Default position
enable_auto_conn	Enable Auto Conn	1
module1_id	First Module	2
module2_id	Second Module	3
module3_id	Third Module	4
automatically_connect_reason	Auto Conn Reason	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed

User Category	Permission
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
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

Chapter 16. 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.

```
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

© Copyright IBM Corp. 2017

User Category	Permission
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 ceritifactes, the maximum already defined

  Troubleshooting: Delete certificate

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

Name	Description	Mandatory	Default
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: base64 -w0 myCert.pfx	Y	N/A

As a security precaution, use the **pki_show_security** command to view the certificate in plain text, and make sure that the certificate text under *Signature Algorithm* does not include the string *MD5*. This will help you avoid a "transcript collision" attack, that can force a hash-construction downgrade to MD5 and reduce expected security. For the vulnerability summary, see the National Vulnerability Database.

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 ceritifactes, the maximum already defined

  Troubleshooting: Delete certificate

Displaying the details of a signed certificate

Use the **pki_show_certificate** command to display the details of a signed certificate.

```
pki_show_certificate name=Name
```

Parameters

Name	Type	Description	Mandatory
name	String	The certificate's symbolic name.	Y

As a security precaution, use this command to view the certificate in plain text, and make sure that the certificate text under *Signature Algorithm* does not include the string *MD5*. This will help you avoid a "transcript collision" attack, that can force a hash-construction downgrade to MD5 and reduce expected security. For the vulnerability summary, see the National Vulnerability Database.

Example:

```
pki_show_certificate name=ibm
```

Field ID	Field output	Default position
certificate	Certificate	1

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- **FAILED_PARSING_CERTIFICATE**
  Failed parsing 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.

Chapter 17. 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.

```
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
pending_ia_cmd	Component Operation	N/A
currently_functioning	Currently Functioning	N/A
requires_service	Requires Service	N/A
service_reason	Service Reason	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

© Copyright IBM Corp. 2017

Listing data counters for the enabled InfiniBand HCA ports

Use the **ib_hca_counter_list** command to list data counters for InfiniBand HCA ports, enabled on modules and flash enclosures.

```
ib_hca_counter_list [ hca_port=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
hca_port	The InfiniBand HCA port to be listed.	N	All IB HCA ports

Example:

```
ib_hca_counter_list
```

Output:

```
Port                        TX Data        RX Data        TX Pkt
--------------------------- -------------- -------------- ------------
1:IB_FlashSystem_Port:4:1   0              0              0
1:IB_FlashSystem_Port:4:3   0              0              0
1:IB_FlashSystem_Port:4:5   0              0              0
1:IB_FlashSystem_Port:4:7   0              0              0
1:IB_Module_Port:12:1       254584779527   252260909954   2055238854
1:IB_Module_Port:12:2       0              0              0
1:IB_Module_Port:13:1       252395242864   254798454598   2061534883
1:IB_Module_Port:13:2       0              0              0
1:IB_Module_Port:8:1        254003578209   254027205845   2055494787
1:IB_Module_Port:8:2        0              0              0

RX Pkt       XmtWait
------------ -----------
0            0
0            0
0            0
0            0
2058771428   103686442
0            0
2060782961   94235849
0            0
2059021166   103031319
0            0
```

Field ID	Field output	Default position
port	Port	1
XmtData	TX Data	2
RcvData	RX Data	3
XmtPkts	TX Pkt	4
RcvPkts	RX Pkt	5
XmtWait	XmtWait	6

Access control

User Category	Permission
Storage administrator	Allowed

User Category	Permission
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing error counters for the enabled InfiniBand HCA ports.

Use the **ib_hca_error_list** command to list error counters for InfiniBand HCA ports, enabled on modules and flash enclosures.

```
ib_hca_error_list [ hca_port=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
hca_port	The InfiniBand HCA port to be listed.	N	All IB HCA ports

Example:

```
ib_hca_error_list
```

Field ID	Field output	Default position
port	Port	1
SymbolErrors	SymErr	2
LinkRecovers	LinkRec	3
LinkDowned	LinkDown	4
RcvErrors	RcvErr	5
RcvRemotePhysErrors	RcvRPErr	6
RcvSwRelayErrors	RcvSRErr	7
XmtDiscards	XmtDisc	8
XmtConstraintErrors	XmtCErr	9
RcvConstraintErrors	RcvCErr	10
LinkIntegrityErrors	LinkIErr	11
ExcBufOverrunErrors	ExcBOEre	12
VL15Dropped	VL15Dr	13

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing the statuses of the enabled InfiniBand HCA ports

Use the **ib_hca_info_list** command to list the statuses of InfiniBand HCA ports, enabled on modules and flash enclosures.

```
ib_hca_info_list [ hca_port=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
hca_port	The InfiniBand HCA port to be listed.	N	All IB HCA ports

Example:

```
ib_hca_info_list
```

Output:

```
Port                        IB Log State   IB Phys State
--------------------------- -------------- ---------------
1:IB_FlashSystem_Port:4:1   NOT SAMPLED    NOT SAMPLED
1:IB_FlashSystem_Port:4:3   NOT SAMPLED    NOT SAMPLED
1:IB_FlashSystem_Port:4:5   NOT SAMPLED    NOT SAMPLED
1:IB_FlashSystem_Port:4:7   NOT SAMPLED    NOT SAMPLED
1:IB_Module_Port:12:1       ACTIVE         LINK UP
1:IB_Module_Port:12:2       INIT           LINK UP
1:IB_Module_Port:13:1       ACTIVE         LINK UP
1:IB_Module_Port:13:2       INIT           LINK UP
1:IB_Module_Port:8:1        ACTIVE         LINK UP
1:IB_Module_Port:8:2        INIT           LINK UP

Link Speed     Link Width
-------------- -------------
NOT SAMPLED    NOT SAMPLED
NOT SAMPLED    NOT SAMPLED
NOT SAMPLED    NOT SAMPLED
NOT SAMPLED    NOT SAMPLED
14.0625 Gbps   X4
14.0625 Gbps   X4
14.0625 Gbps   X4
14.0625 Gbps   X4
14.0625 Gbps   X4
14.0625 Gbps   X4
```

Field ID	Field output	Default position
port	Port	1
log_state	IB Log State	2
phys_state	IB Phys State	3
link_speed	Link Speed	4
link_width	Link Width	5
link_width_sup	Link Width Sup	N/A
link_speed_sup	Link Speed Sup	N/A
link_speed_enabled	Link Speed Ena	N/A
link_width_enabled	Link Width Ena	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing the configured InfiniBand switches

Use the **switch_list** command to list the configured InfiniBand switches.

```
switch_list [ switch=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
switch	IB switch to list.	N	All IB switches

Example:

```
switch_list
```

Output:

```
Switch          Switch GUID        Status   MGMT OK   Ports OK   Power OK   BBU OK
--------------- ------------------ -------- --------- ---------- ---------- --------
1:IB_Switch:1   E41D2D03003C9900   OK       yes       yes        yes        yes
1:IB_Switch:2   E41D2D03003C9A80   OK       yes       yes        yes        yes

Cont.:

Fan OK   Temp OK   Volt OK   Boot Time             FW         Serial
-------- --------- --------- -------------------- ---------- --------------
yes      yes       yes       04/04/2016 10:51:18   3.5.0500   MT1523X09088
yes      yes       yes       27/03/2016 15:28:31   3.5.0500   MT1523X09091
```

Field ID	Field output	Default position
component_id	Switch	1
status	Status	3
sw_mgmt_status	MGMT Status	N/A
num_of_down_ports	Down Ports	N/A
mgmt_ok	MGMT OK	4
ports_ok	Ports OK	5
power_ok	Power OK	6
bbu_ok	BBU OK	7
fan_ok	Fan OK	8
temp_ok	Temp OK	9
volt_ok	Volt OK	10
fw	FW	12

Field ID	Field output	Default position
mgmt_serial_number	Serial	13
mgmt_part_number	Part No	N/A
mgmt_asic_rev	ASIC Rev	N/A
mgmt_hw_rev	HW Rev	N/A
chassis_serial_number	Chassis Serial	N/A
chassis_part_number	Chassis Part No	N/A
chassis_asic_rev	Chassis ASIC Rev	N/A
chassis_hw_rev	Chassis HW Rev	N/A
original_mgmt_serial_number	Original Serial	N/A
original_mgmt_part_number	Original Part No	N/A
original_mgmt_asic_rev	Original ASIC Rev	N/A
original_mgmt_hw_rev	Original HW Rev	N/A
original_chassis_serial_number	Original Chassis Serial	N/A
original_chassis_part_number	Original Chassis Part No	N/A
original_chassis_asic_rev	Original Chassis ASIC Rev	N/A
original_chassis_hw_rev	Original Chassis HW Rev	N/A
currently_functioning	Currently Functioning	N/A
mgmt_guid	Managememt GUID	N/A
requires_service	Requires Service	N/A
service_reason	Service Reason	N/A
fabric_switch_info.miss_count	Miss Count	N/A
fabric_switch_info.num_of_ports	Ports	N/A
fabric_switch_info.image_guid	FW GUID	N/A
fabric_switch_info.node_guid	Switch GUID	2
fabric_switch_info.dev_id	Device ID	N/A
fabric_switch_info.dev_rev	Device Revision	N/A
fabric_switch_info.vendor_id	Vendor ID	N/A
fabric_switch_info.name	Name	N/A
fabric_switch_info.mlx_dev_id	Ext Device ID	N/A
fabric_switch_info.mlx_hw_rev	Ext Device Revision	N/A
fabric_switch_info.boot_time	Boot Time	11
fabric_switch_info.fw_build_id	FW BUILD ID	N/A
fabric_switch_info.fw_rev	FW Ver	N/A
fabric_switch_info.fw_build_date	FW Build Date	N/A
fabric_switch_info.psid	PSID	N/A
used_power	Total Power Used	N/A
power_capacity	Total Power Capacity	N/A
power_available	Total Power Available	N/A
projected_max_used_power	Projected Max User Power	N/A
bbu_runtime	Battery Runtime	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing the configured InfiniBand switch management addresses

Use the **switch_mgmt_ip_list** command to list the configured InfiniBand switch management addresses.

```
switch_mgmt_ip_list [ switch=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
switch	IB switch to list.	N	All IB switches

Example:

```
switch_mgmt_ip_list
```

Output:

```
Switch          Type     NAT IP        Router          Real IP         Status
--------------- -------- ------------- --------------- --------------- --------
1:IB_Switch:1   IPOIB    14.10.255.1   1:IB_Switch:1   14.10.255.1     OK
1:IB_Switch:1   MGMT1    14.10.254.1   1:Module:1      192.168.0.254   OK
1:IB_Switch:1   MGMT2    14.10.253.1   1:Module:4      192.168.1.254   OK
1:IB_Switch:1   SERIAL   14.10.10.3    1:Module:3      14.10.10.3      OK
1:IB_Switch:2   IPOIB    14.10.255.2   1:IB_Switch:2   14.10.255.2     OK
1:IB_Switch:2   MGMT1    14.10.254.2   1:Module:2      192.168.0.254   OK
1:IB_Switch:2   MGMT2    14.10.253.2   1:Module:3      192.168.1.254   OK
1:IB_Switch:2   SERIAL   14.10.10.4    1:Module:4      14.10.10.4      OK
```

Field ID	Field output	Default position
switch_id	Switch	1
type	Type	2
nat_ip	NAT IP	3
router	Router	4
real_ip	Real IP	5
status	Status	6

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed

User Category	Permission
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing the configured InfiniBand switch firmware versions

Use the **switch_fw_list** command to list the configured InfiniBand switch firmware versions.

```
switch_fw_list [ switch=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
switch	IB switch to list.	N	All IB switches

Example:

```
switch_fw_list
```

Output:

```
Switch          Type              Version
--------------- ----------------- ----------
1:IB_Switch:1   ASIC              9.3.7170
1:IB_Switch:1   BBU1              703
1:IB_Switch:1   BBU2              703
1:IB_Switch:1   BIOS              4.6.5
1:IB_Switch:1   CPLD_PORT1        4
1:IB_Switch:1   CPLD_SWITCH_BRD   7
1:IB_Switch:1   CPLD_TOR          9
1:IB_Switch:1   MGMT              3.5.0500
1:IB_Switch:1   PSU1              404
1:IB_Switch:1   PSU2              404
1:IB_Switch:2   ASIC              9.3.7170
1:IB_Switch:2   BBU1              703
1:IB_Switch:2   BBU2              703
1:IB_Switch:2   BIOS              4.6.5
1:IB_Switch:2   CPLD_PORT1        4
1:IB_Switch:2   CPLD_SWITCH_BRD   7
1:IB_Switch:2   CPLD_TOR          9
1:IB_Switch:2   MGMT              3.5.0500
1:IB_Switch:2   PSU1              404
1:IB_Switch:2   PSU2              404
```

Field ID	Field output	Default position
switch_id	Switch	1
type	Type	2
version	Version	3
original_version	Original Version	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing the configured InfiniBand switch power values

Use the **switch_power_list** command to list the configured InfiniBand switch power values for PSUs and BBUs.

```
switch_power_list [ switch=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
switch	IB switch to list.	N	All IB switches

Example:

```
switch_power_list
```

Output:

```
Switch          Type   Power W   Voltage V   Current A   Capacity W   Feed   Status
--------------- ------ --------- ----------- ----------- ------------ ------ --------
1:IB_Switch:1   BBU1   N/A       N/A         N/A         330.00       NA     OK
1:IB_Switch:1   BBU2   N/A       N/A         N/A         330.00       NA     OK
1:IB_Switch:1   PS1    46.00     12.11       2.56        400.00       AC     OK
1:IB_Switch:1   PS2    44.00     12.19       2.75        400.00       AC     OK
1:IB_Switch:2   BBU1   N/A       N/A         N/A         330.00       NA     OK
1:IB_Switch:2   BBU2   N/A       N/A         N/A         330.00       NA     OK
1:IB_Switch:2   PS1    47.00     12.05       2.56        400.00       AC     OK
1:IB_Switch:2   PS2    44.00     12.14       2.81        400.00       AC     OK
```

Field ID	Field output	Default position
switch_id	Switch	1
type	Type	2
power	Power W	3
voltage	Voltage V	4
current	Current A	5
capacity	Capacity W	6
feed	Feed	7
status	Status	8

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing the configured InfiniBand switch voltage values

Use the **switch_voltage_list** command to list the configured InfiniBand switch voltage values.

```
switch_voltage_list [ switch=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
switch	IB switch to list.	N	All IB switches

Example:

```
switch_voltage_list
```

Output:

```
Switch          Type            Expected   Actual   Status   High    Low
--------------- --------------- ---------- -------- -------- ------- -------
1:IB_Switch:1   1.05V LAN       1.50       1.49     OK       1.72    1.27
1:IB_Switch:1   Asic 1.2V       1.20       1.20     OK       1.38    1.02
1:IB_Switch:1   Asic 1.8V       1.80       1.81     OK       2.06    1.53
1:IB_Switch:1   Asic 3.3V       3.30       3.31     OK       3.79    2.80
1:IB_Switch:1   BBU1            12.00      12.50    OK       13.80   10.19
1:IB_Switch:1   BBU2            12.00      12.50    OK       13.80   10.19
1:IB_Switch:1   CPU 0.9V        0.90       0.85     OK       1.03    0.77
1:IB_Switch:1   CPU 1.05V       1.05       1.03     OK       1.21    0.89
1:IB_Switch:1   CPU 1.8V        1.80       1.78     OK       2.06    1.53
1:IB_Switch:1   CPU/PCH 1.05V   1.05       1.00     OK       1.10    0.81
1:IB_Switch:1   DDR3 0.675V     0.68       0.66     OK       0.78    0.56
1:IB_Switch:1   DDR3 1.35V      1.35       1.34     OK       1.55    1.14
1:IB_Switch:1   PS1 vout 12V    12.00      12.11    OK       13.80   10.19
1:IB_Switch:1   PS2 vout 12V    12.00      12.19    OK       13.80   10.19
1:IB_Switch:1   SYS 3.3V        3.30       3.31     OK       3.79    2.80
1:IB_Switch:1   USB 5V          5.00       5.01     OK       5.75    4.25
1:IB_Switch:1   Vcore SX        0.95       0.96     OK       1.09    0.81
1:IB_Switch:2   1.05V LAN       1.50       1.52     OK       1.72    1.27
1:IB_Switch:2   Asic 1.2V       1.20       1.21     OK       1.38    1.02
1:IB_Switch:2   Asic 1.8V       1.80       1.81     OK       2.06    1.53
1:IB_Switch:2   Asic 3.3V       3.30       3.32     OK       3.79    2.80
1:IB_Switch:2   BBU1            12.00      12.50    OK       13.80   10.19
1:IB_Switch:2   BBU2            12.00      12.50    OK       13.80   10.19
1:IB_Switch:2   CPU 0.9V        0.90       0.86     OK       1.03    0.77
1:IB_Switch:2   CPU 1.05V       1.05       1.06     OK       1.21    0.89
1:IB_Switch:2   CPU 1.8V        1.80       1.83     OK       2.06    1.53
1:IB_Switch:2   CPU/PCH 1.05V   1.05       1.02     OK       1.10    0.81
1:IB_Switch:2   DDR3 0.675V     0.68       0.68     OK       0.78    0.56
1:IB_Switch:2   DDR3 1.35V      1.35       1.37     OK       1.55    1.14
1:IB_Switch:2   PS1 vout 12V    12.00      12.05    OK       13.80   10.19
1:IB_Switch:2   PS2 vout 12V    12.00      12.14    OK       13.80   10.19
1:IB_Switch:2   SYS 3.3V        3.30       3.41     OK       3.79    2.80
1:IB_Switch:2   USB 5V          5.00       5.16     OK       5.75    4.25
1:IB_Switch:2   Vcore SX        0.95       0.96     OK       1.09    0.81
```

Field ID	Field output	Default position
switch_id	Switch	1
type	Type	2
expected	Expected	3
actual	Actual	4
status	Status	5
high_margin	High	6
low_margin	Low	7

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing the configured InfiniBand switch temperature values

Use the **switch_temp_list** command to list the configured InfiniBand switch temperature values.

```
switch_temp_list [ switch=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
switch	IB switch to list.	N	All IB switches

Example:

```
switch_temp_list
```

Output:

```
Switch          Type         Actual   Alert    Critical   Status
--------------- ------------ -------- -------- ---------- --------
1:IB_Switch:1   BBU1         19.60    60.00    60.00      OK
1:IB_Switch:1   BBU2         19.50    60.00    60.00      OK
1:IB_Switch:1   MGMT_AMB     20.50    120.00   120.00     OK
1:IB_Switch:1   MGMT_CPU     25.00    120.00   120.00     OK
1:IB_Switch:1   MGMT_CPU1    22.00    120.00   120.00     OK
1:IB_Switch:1   MGMT_CPU2    25.00    120.00   120.00     OK
1:IB_Switch:1   MGMT_PORTS   22.00    120.00   120.00     OK
1:IB_Switch:1   MGMT_SX      28.00    105.00   110.00     OK
1:IB_Switch:1   PS1          24.00    120.00   120.00     OK
1:IB_Switch:1   PS2          24.00    120.00   120.00     OK
1:IB_Switch:2   BBU1         20.00    60.00    60.00      OK
1:IB_Switch:2   BBU2         19.89    60.00    60.00      OK
1:IB_Switch:2   MGMT_AMB     21.00    120.00   120.00     OK
1:IB_Switch:2   MGMT_CPU     26.00    120.00   120.00     OK
1:IB_Switch:2   MGMT_CPU1    25.00    120.00   120.00     OK
1:IB_Switch:2   MGMT_CPU2    19.00    120.00   120.00     OK
1:IB_Switch:2   MGMT_PORTS   22.50    120.00   120.00     OK
1:IB_Switch:2   MGMT_SX      28.00    105.00   110.00     OK
1:IB_Switch:2   PS1          24.00    120.00   120.00     OK
1:IB_Switch:2   PS2          25.00    120.00   120.00     OK
```

Field ID	Field output	Default position
switch_id	Switch	1
type	Type	2
actual	Actual	3
alert	Alert	4
critical	Critical	5
status	Status	6

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed

User Category	Permission
Read-only users	Allowed
Technicians	Allowed

Listing the configured InfiniBand switch fan parts

Use the **switch_fan_part_list** command to list the configured InfiniBand switch fan parts.

```
switch_fan_part_list [ switch=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
switch	IB switch to list.	N	All IB switches

Example:

```
switch_fan_part_list
```

Output:

```
Switch          Type      Part No         Serial No      HW Rev   Speed      Status
--------------- --------- --------------- -------------- -------- ---------- --------
1:IB_Switch:1   FAN1-F1   MTEF-FANF-B     MT1523X09083   A2       10608.00   OK
1:IB_Switch:1   FAN1-F2   MTEF-FANF-B     MT1523X09083   A2       8998.00    OK
1:IB_Switch:1   FAN2-F1   MTEF-FANF-B     MT1523X09075   A2       10526.00   OK
1:IB_Switch:1   FAN2-F2   MTEF-FANF-B     MT1523X09075   A2       8939.00    OK
1:IB_Switch:1   FAN3-F1   MTEF-FANF-B     MT1523X09082   A2       10608.00   OK
1:IB_Switch:1   FAN3-F2   MTEF-FANF-B     MT1523X09082   A2       8998.00    OK
1:IB_Switch:1   FAN4-F1   MTEF-FANF-B     MT1523X09077   A2       10691.00   OK
1:IB_Switch:1   FAN4-F2   MTEF-FANF-B     MT1523X09077   A2       9242.00    OK
1:IB_Switch:1   PSU1      MTEF-PSF-AC-B   MT1523X09040   A5       14464.00   OK
1:IB_Switch:1   PSU2      MTEF-PSF-AC-B   MT1523X09044   A5       14592.00   OK
1:IB_Switch:2   FAN1-F1   MTEF-FANF-B     MT1523X09065   A2       10526.00   OK
1:IB_Switch:2   FAN1-F2   MTEF-FANF-B     MT1523X09065   A2       9118.00    OK
1:IB_Switch:2   FAN2-F1   MTEF-FANF-B     MT1523X09072   A2       10445.00   OK
1:IB_Switch:2   FAN2-F2   MTEF-FANF-B     MT1523X09072   A2       9118.00    OK
1:IB_Switch:2   FAN3-F1   MTEF-FANF-B     MT1523X09062   A2       10445.00   OK
1:IB_Switch:2   FAN3-F2   MTEF-FANF-B     MT1523X09062   A2       8998.00    OK
1:IB_Switch:2   FAN4-F1   MTEF-FANF-B     MT1523X09073   A2       10526.00   OK
1:IB_Switch:2   FAN4-F2   MTEF-FANF-B     MT1523X09073   A2       8998.00    OK
1:IB_Switch:2   PSU1      MTEF-PSF-AC-B   MT1523X09035   A5       14464.00   OK
1:IB_Switch:2   PSU2      MTEF-PSF-AC-B   MT1523X09033   A5       14464.00   OK
```

Field ID	Field output	Default position
switch_id	Switch	1
type	Type	2
part_no	Part No	3
serial_no	Serial No	4
hw_rev	HW Rev	5
speed	Speed	6
status	Status	7

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing the configured InfiniBand switch PSUs

Use the **switch_psu_list** command to list the configured InfiniBand switch PSUs.

```
switch_psu_list [ switch_psu=ComponentId ] [ switch=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
switch_psu	IB switch PSU to list.	N	All IB switch PSUs
switch	IB switch whose BBUs are to be listed.	N	All IB switches

Example:

```
switch_psu_list
```

Output:

```
Component ID          Status   Sensor Status   Power W   Voltage V
--------------------- -------- --------------- --------- -----------
1:IB_Switch_PSU:1:1   OK       OK              46.00     12.11
1:IB_Switch_PSU:1:2   OK       OK              44.00     12.19
1:IB_Switch_PSU:2:1   OK       OK              47.00     12.05
1:IB_Switch_PSU:2:2   OK       OK              44.00     12.14

Cont.:

Current A   Capacity W
----------- ------------
2.94        400.00
2.50        400.00
2.69        400.00
3.25        400.00
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
sensor_status	Sensor Status	3
power	Power W	4
voltage	Voltage V	5
current	Current A	6
capacity	Capacity W	7
temperature	Temperature	N/A

Field ID	Field output	Default position
fan_speed	Fan Speed	N/A
fan_sensor_status	Fan Status	N/A
serial_number	Serial No	N/A
original_serial_number	Original Serial No	N/A
part_number	Part No	N/A
original_part_number	Original Part No	N/A
hw_rev	HW Rev	N/A
requires_service	Requires Service	N/A
service_reason	Service Reason	N/A
currently_functioning	Currently Functioning	N/A
switch_id	Switch ID	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing the configured InfiniBand switch BBUs

Use the **switch_bbu_list** command to list the configured InfiniBand switch BBUs.

```
switch_bbu_list [ switch_bbu=ComponentId ] [ switch=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
switch_bbu	IB switch BBU to list.	N	All IB switch BBUs
switch	IB switch to list the BBUs of.	N	All IB switches

Example:

```
switch_bbu_list
```

Output:

```
Component ID          Status   Sensor Status   Remaining capacity   Full charged capacity
--------------------- -------- --------------- -------------------- ----------------------
1:IB_Switch_BBU:1:1   OK       OK              71520 mWh            80660 mWh
1:IB_Switch_BBU:1:2   OK       OK              46060 mWh            55520 mWh
1:IB_Switch_BBU:2:1   OK       OK              80800 mWh            80800 mWh
1:IB_Switch_BBU:2:2   OK       OK              77840 mWh            80120 mWh

Cont.:

Percent Charged   Charger Status   Calibration Status   Last Recondition Date
----------------- ---------------- -------------------- ----------------------
100%              Fully charged    N/A                  N/A
100%              Fully charged    N/A                  N/A
100%              Fully charged    N/A                  N/A
100%              Fully charged    N/A                  N/A
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
sensor_status	Sensor Status	3
relative_capacity	Remaining capacity	4
absolute_capacity	Full charged capacity	5
relative_capacity_percent	Percent Charged	6
charging_state	Charger Status	7
calibration_status	Calibration Status	8
last_calibration_date	Last Recondition Date	9
fw_version	FW	N/A
serial_number	Serial No	N/A
original_serial_number	Original Serial No	N/A
part_number	Part No	N/A
original_part_number	Original Part No	N/A
voltage	Voltage	N/A
exp_voltage	Expected Voltage	N/A
temperature	Temperature	N/A
manufacture_date	Manufacture Date	N/A
designed_capacity	Designed Capacity	N/A
absolute_charge	Absolute Charge	N/A
test_status	Test Status	N/A
last_test_date	Last Test Date	N/A
can_charge	Charge	N/A
can_discharge	Discharge	N/A
charge_rate	Charge Rate	N/A
max_error	Max Error	N/A
hw_rev	HW Rev	N/A
requires_service	Requires Service	N/A
service_reason	Service Reason	N/A
currently_functioning	Currently Functioning	N/A
switch_id	Switch ID	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing the configured InfiniBand switch fans

Use the **switch_fan_list** command to list the configured InfiniBand switch fans.

```
switch_fan_list [ switch_fan=ComponentId ] [ switch=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
switch_fan	IB switch fan to list.	N	All IB switch fans
switch	IB switch whose BBUs are to be listed.	N	All IB switches

Example:

```
switch_fan_list
```

Output:

```
Component ID          Status   Speed      Sensor Status   Peer Speed   Peer Sensor Status
--------------------- -------- ---------- --------------- ------------ --------------------
1:IB_Switch_Fan:1:1   OK       10608.00   OK              8998.00      OK
1:IB_Switch_Fan:1:2   OK       10526.00   OK              8939.00      OK
1:IB_Switch_Fan:1:3   OK       10608.00   OK              8998.00      OK
1:IB_Switch_Fan:1:4   OK       10691.00   OK              9242.00      OK
1:IB_Switch_Fan:2:1   OK       10526.00   OK              9118.00      OK
1:IB_Switch_Fan:2:2   OK       10445.00   OK              9118.00      OK
1:IB_Switch_Fan:2:3   OK       10445.00   OK              8998.00      OK
1:IB_Switch_Fan:2:4   OK       10526.00   OK              8998.00      OK
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
speed	Speed	3
sensor_status	Sensor Status	4
peer_speed	Peer Speed	5
peer_sensor_status	Peer Sensor Status	6
serial_number	Serial No	N/A
original_serial_number	Original Serial No	N/A
part_number	Part No	N/A
original_part_number	Original Part No	N/A
hw_rev	HW Rev	N/A

Field ID	Field output	Default position
requires_service	Requires Service	N/A
service_reason	Service Reason	N/A
currently_functioning	Currently Functioning	N/A
switch_id	Switch ID	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Chapter 18. 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.

```
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

© Copyright IBM Corp. 2017

- **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
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

Name	Type	Description	Mandatory	Default
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
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_attrib
storage_admin_role
read_only_role
session_cache_period    20
bind_time_limit         20
user_id_attrib          objectSiD
first_expiration_event  30
second_expiration_event 14
third_expiration_event  7
use_ssl                 no
xiv_user
```

Field ID	Field output	Default position
name	Name	1
value	Value	2

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Configuring LDAP in the system

Use the **ldap_config_set** command to configure general system parameters that control user authentication against LDAP servers.

```
ldap_config_set [ user_name_attrib=LdapAttrib ] [ xiv_group_attrib=LdapAttrib ]
 [ storage_admin_role=LdapRole ] [ read_only_role=LdapRole ]
 [ security_admin_role=LdapRole ] [ storage_integration_admin_role=LdapRole ]
 [ xiv_host_profiler_role=LdapRole ] [ use_ssl=<yes|no> ] [ user_id_attrib=LdapAttrib ]
 [ session_cache_period=Minutes ] [ bind_time_limit=Seconds ]
 [ first_expiration_event=Days ] [ second_expiration_event=Days ]
 [ third_expiration_event=Days ] [ version=LdapVersion ] [ xiv_user=LdapAttrib ]
 [ xiv_password=LdapAttrib ]
 [ server_type=<SUN DIRECTORY|MICROSOFT ACTIVE DIRECTORY|OPEN LDAP> ]
 [ group_search_depth=Depth ] [ group_search_max_queries=Number ]
 [ group_search_stop_when_found=<yes|no> ]
```

Parameters

Name	Type	Description	Mandatory	Default
user_name_attrib	String	User name attribute for queries. If not specified, it is set to `uid` for SUN Directory servers and `userPrincipalName` for Microsoft Active Directory servers.	N	According to server type
xiv_group_attrib	String	LDAP attribute designated to hold system-mapped roles.	N	none
storage_admin_role	String	LDAP value mapped to the Storage Administrator role. Multiple (up to 8) values are supported and must be separated with a semicolon (;). Multiple roles are not available for SUN Directory LDAP Servers.	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

Name	Type	Description	Mandatory	Default
xiv_host_ profiler_role	String	LDAP value mapped to the XIV Host Profiler role. Multiple (up to 8) values are supported and must be separated with a semicolon (;).	N	none
use_ssl	Boolean	Indicates whether secure LDAP is mandatory.	N	no
user_id_attrib	String	The LDAP attribute set to identify the user (in addition to user DN) when recording user operations in the event log.	N	objectSiD
session_cache_ period	Positive integer	Duration of keeping user credentials before attempting to re-login the user.	N	20
bind_time_limit	Positive integer	The duration after which the next LDAP server on the LDAP server list will be called.	N	0. If set to the default, the LDAP server is contacted for every command. Performance issues depend on its availability.
first_ expiration_ event	Positive integer	The number of days before the expiration of certificate, when the first alert is issued (severity: warning).	N	30/14/7 (third is smallest)
second_ expiration_ event	Positive integer	The number of days before the expiration of certificate, when the second alert is issued (severity: warning).	N	30/14/7 (third is smallest)
third_ expiration_ event	Positive integer	The number of days before the expiration of certificate, when the third alert is issued (severity: warning).	N	30/14/7 (third is smallest)
version	Positive integer	Version of LDAP used (only version 3 is supported).	N	3
xiv_user	String	The user for LDAP queries.	N	none
xiv_password	String	The password of user for LDAP queries.	N	none

Name	Type	Description	Mandatory	Default
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 storage_admin_role="CN=EXMPLEAdmins,OU=EXMPLELab,DC=CROSSQA,DC=COM"
read_only_role="CN=EXMPLEOpers,OU=EXMPLELab,DC=CROSSQA,DC=COM"
user_id_attrib=sAMAccountName user_name_attrib=sAMAccountName
xiv_group_attrib=memberOf xiv_user="CN=allmighty,CN=Users,DC=CROSSQA,DC=COM"
xiv_password=the_password
security_admin_role="CN=EXMPLESecAdmins,OU=EXMPLELab,DC=CROSSQA,DC=COM"
```

Output:

```
Command 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

- **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 shold 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 shold 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 shold be between 0 and 8
- **LDAP_XIV_HOST_PROFILER_ROLE_HAS_WRONG_NUMBER_OF_PARTS**

  LDAP XIV Host Profiler 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
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

User Category	Permission
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:: AQUAAAAAAAUVAAAAYcKhSnhmt01IPSuAbQQAAA==
stdout        3
stdout        4
```

Field ID	Field output	Default position
name	Name	1
index	Index	2
value	Value	3

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- **LOGIN_FAILURE_USER_NOT_AUTHENTICATED_BY_LDAP_SERVER**
  User *User Name* was not authenticated by LDAP server '*Server FQDN*'.
- **LOGIN_FAILURE_USER_MISSING_GROUP_ATTRIBUTE**
  User *User Name* is missing the group attribute '*Attribute*'.
- **LDAP_SERVER_NOT_FOUND**
  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. Any symbols are allowed, except the following: • double quotation (") • single quotation or apostrophe (') • grave accent (`) Passwords are case sensitive.	Y	N/A
password_verify	N/A	Password verification, which must be equal to the value of `password`.	Y	N/A
category	Enumeration	The role of the user to be created. Available roles: storageadmin, applicationadmin, operationsadmin, securityadmin, readonly and opsadmin.	Y	N/A

Name	Type	Description	Mandatory	Default
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
user_group	Object name	User group to be deleted.	Y

A user group can be deleted, even when it is associated with hosts or clusters. It can be deleted while in LDAP Authentication mode.

A user group can be deleted, even when it contains users. Deleting the user group does not delete the users contained in this group.

Example:

```
user_group_delete user_group=ug1
```

Output:

```
Command 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
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**
  Designtated 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 must have between 6 and 12 characters. Any symbols are allowed, except the following: • double quotation (") • single quotation or apostrophe (') • grave accent (`) Passwords are case sensitive.	N	Retains the current password.
password_verify	N/A	Verification of the password: Must be equal to the password.	N	Retains the current password.
email_address	N/A	Email address of 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.

Name	Type	Description	Mandatory	Default
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.

A user with the predefined password `admin` can change the passwords of other users. 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**

  Designtated 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 [ size=GB ] [ max_pools=MaxPools ]
 [ max_volumes=MaxVolumes ] [ max_cgs=MaxCGs ] [ max_mirrors=MaxMirrors ]
 [ max_dms=MaxDataMigrations ] [ perf_class=perfClassName ] [ ldap_id=LdapRole ]
```

Parameters

Name	Type	Description	Mandatory	Default
domain	Object name	The name of the domain to be created.	Y	N/A

Name	Type	Description	Mandatory	Default
size	Integer	Defines the sum of the sizes of all the pools associated with the domain, in gigabytes.	N	0
max_pools	Positive integer	The maximum number of pools that can be associated with this domain.	N	0
max_volumes	Positive integer	The maximum number of volumes that can be associated with all the pools in this domain.	N	0
max_cgs	Integer	The maximum number of consistency groups that can be associated with this domain.	N	512
max_mirrors	Positive integer	The maximum number of mirrors that can be associated with this domain.	N	0
max_dms	Positive integer	The maximum number of data migrations that can be associated with this domain.	N	0
perf_class	Object name	Name of a performance class.	N	none
ldap_id	String	The name to be associated with this domain in LDAP.	N	The domain name

Example:

```
domain_create domain=d1 size=1000
```

Output:

```
Command 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?
- **DOMAIN_SIZE_SMALL**

  Domain size is smaller than minimal pool size. Are you sure?

Return codes

- **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
- **PERF_CLASS_ASSOCIATED_WITH_HOSTS**

  Performance Class *Performance Class* is already in use by host.
- **PERF_CLASS_ASSOCIATED_WITH_VOLUMES**

  Performance Class *Performance Class* is already in use by volume.
- **DOMAIN_INSUFFICIENT_CAPACITY**

  There is not enough capacity available for the domain.
- **DOMAIN_INSUFFICIENT_VOLUMES**

  There are not enough volumes available for the domain.
- **DOMAIN_INSUFFICIENT_POOLS**

  There are not enough pools available for the domain.
- **DOMAIN_INSUFFICIENT_CGS**

  There are not enough cons groups available for the domain.
- **DOMAIN_INSUFFICIENT_DMS**

  There are not enough data migration available for the domain.

Updating a domain definition

Use the **domain_update** command to update a domain definition.

```
domain_update domain=DomainName [ size=GB ] [ max_pools=MaxPools ]
 [ max_volumes=MaxVolumes ] [ max_cgs=MaxCGs ] [ max_mirrors=MaxMirrors ]
 [ max_dms=MaxDataMigrations ] [ perf_class=perfClassName ] [ ldap_id=LdapRole ]
```

Parameters

Name	Type	Description	Mandatory	Default
domain	Object name	The name of the domain to be updated.	Y	N/A
size	Integer	Defines the sum of the sizes of all the pools associated with the domain, in gigabytes.	N	Current value.

Name	Type	Description	Mandatory	Default
max_pools	Positive integer	The maximum number of pools that can be associated with this domain.	N	Current value.
max_volumes	Positive integer	The maximum number of volumes that can be associated with all the pools in this domain.	N	Current value.
max_cgs	Integer	The maximum number of consistency groups that can be associated with this domain.	N	Current value.
max_mirrors	Positive integer	The maximum number of mirrors that can be associated with this domain.	N	Current value.
max_dms	Positive integer	The maximum number of data migrations that can be associated with this domain.	N	Current value.
perf_class	Object name	Name of a performance class.	N	Current value.
ldap_id	String	The name to be associated with this domain in LDAP.	N	Current value.

Example:

```
domain_update domain=d1 size=10000 max_pools=5 max_volumes=100
```

Output:

```
Command 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?

- **DOMAIN_SIZE_SMALL**

  Domain size is smaller than minimal pool size. Are you sure?

Return codes

- **DOMAIN_DOESNT_EXIST**

  Domain does not exist.
- **PERF_CLASS_BAD_NAME**

  Performance Class does not exist
- **PERF_CLASS_ASSOCIATED_WITH_HOSTS**

  Performance Class *Performance Class* is already in use by host.
- **PERF_CLASS_ASSOCIATED_WITH_VOLUMES**

  Performance Class *Performance Class* is already in use by volume.
- **DOMAIN_SIZE_TOO_SMALL**

  Domain usage exceeds requested size.
- **DOMAIN_MAX_VOLUMES_TOO_SMALL**

  Domain already has more than the specified number of volumes.
- **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_INSUFFICIENT_CAPACITY**

  There is not enough capacity available for the domain.
- **DOMAIN_INSUFFICIENT_VOLUMES**

  There are not enough volumes available for the domain.
- **DOMAIN_INSUFFICIENT_POOLS**

  There are not enough pools available for the domain.
- **DOMAIN_INSUFFICIENT_CGS**

  There are not enough cons groups available for the domain.
- **DOMAIN_INSUFFICIENT_DMS**

  There are not enough data migration available for the domain.

Renaming a domain

Use the **domain_rename** command to rename a domain.

```
domain_rename domain=DomainName new_name=Name
```

Parameters

Name	Type	Description	Mandatory
new_name	Object name	Name of the domain.	Y
domain	Object name	New name of the domain.	Y

Example:

```
domain_rename domain=domain1 new_name=domain2
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- **DOMAIN_ALREADY_EXISTS**
  A domain with this name already exists.
- **DOMAIN_DOESNT_EXIST**
  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

User Category	Permission
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_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="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="50713d00000"/>
 <perf_class value="QoS1"/>>
 <dn value="Domain1"/>
</domain>
```

Field ID	Field output	Default position
name	Name	1
ldap_id	LDAP ID	2
size	Size	3
size_MiB	Size (MiB)	N/A
total_pool_size	Total Pools (GB)	4
total_pool_size_MiB	Total Pools (MiB)	N/A
empty_space	Empty (GB)	5
empty_space_MiB	Empty (MiB)	N/A
max_pools	Max Pools	6
used_pools	Pools	7
max_volumes	Max Volumes	8
used_volumes	Volumes	9
max_mirrors	Max Mirrors	10
used_mirrors	Mirrors	11
max_dms	Max Data Migrations	12
used_dms	Data Migrations	13
max_cgs	Max CGs	14
used_cgs	CGs	15
perf_class	Performance Class	16
managed	Managed	17
id	ID	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Allowed
Read-only users	Allowed
Technicians	Disallowed

Listing users per domain

Use the **domain_list_users** command to list users associated with domain(s).

```
domain_list_users  [ domain=DomainName ] [ user=UserName ] [ category=Category ]
[ show_users=<all|active> ]
```

Parameters

Name	Type	Description	Mandatory	Default
domain	Object name	Name of a domain.	N	All domains.
user	Object name	Name of a user.	N	All users.
category	Enumeration	The roles of the users to be listed. Available options are: storageadmin, readonly, applicationadmin and storageintegrationadmin.	N	All categories.
show_users	Enumeration	Indicates whether to list all internal users, or only active internal users.	N	active

Example:

```
domain_list_users domain=d1
```

Output:

```
Domain   User        Category
-------- ----------- --------------
d1       d1_domain   storageadmin
```

Field ID	Field output	Default position
domain_name	Domain	1
user_name	User	2
category	Category	3

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Allowed
Read-only users	Allowed
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
size	Size	3
size_MiB	Size (MiB)	N/A
total_pool_size	Total Pools (GB)	4
total_pool_size_MiB	Total Pools (MiB)	N/A
empty_space	Empty (GB)	5
empty_space_MiB	Empty (MiB)	N/A
max_pools	Max Pools	6
used_pools	Pools	7
max_volumes	Max Volumes	8
used_volumes	Volumes	9
max_mirrors	Max Mirrors	10
used_mirrors	Mirrors	11
max_dms	Max Data Migrations	12
used_dms	Data Migrations	13
max_cgs	Max CGs	14
used_cgs	CGs	15
perf_class	Performance Class	16
managed	Managed	17

Field ID	Field output	Default position
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 by included in the destination domains.
- **DESTINATION_IS_NOT_IN_DESTGROUP_DOMAINS**
  Destinations must by included in the destination group domains.
- **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.
- **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

Name	Type	Description	Mandatory
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 by included in the destination domains.
- **DESTINATION_IS_NOT_IN_RULE_DOMAINS**

  Destination must by included in the rule domains.
- **DESTINATION_IS_NOT_IN_DESTGROUP_DOMAINS**

  Destinations must by 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 by included in the rule domains.
- **ESCALATION_RULE_NOT_IN_RULE_DOMAINS**

  Escalation rule must belong to rule domains.
- **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**

Designtated 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
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
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed

User Category	Permission
Technicians	Disallowed

Return codes

- **POOL_DOES_NOT_EXIST**

  Storage Pool does not exist
- **NO_FREE_CAPACITY_IN_DOMAIN**

  There is not enough free 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_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.
- **MAPPED_HOSTS_NOT_IN_DESTINATION**

  A host that is mapped to a volume in the pool is not associated with the target domain.
- **MAPPED_CLUSTERS_NOT_IN_DESTINATION**

  A cluster that is mapped to a volume in the pool is not associated with the target domain.
- **NO_SPACE**

  The system does not have enough free space for the requested Storage Pool size

Removing a pool from a domain

Use the **domain_remove_pool** command to remove a pool from a domain.

```
domain_remove_pool domain=DomainName pool=PoolName [ adjust=<yes|no> ]
```

Parameters

Name	Type	Description	Mandatory	Default
domain	Object name	The name of the domain.	Y	N/A
pool	Object name	The pool name.	Y	N/A

Name	Type	Description	Mandatory	Default
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_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
- **NO_SPACE**
  The system does not have enough free space for the requested Storage Pool size

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.
- **NO_FREE_CAPACITY_IN_DOMAIN**

  There is not enough free space in the 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_DMS_REACHED**

  The domain exceeds the maximum allowed number of data migrations.
- **DOMAIN_MAX_CONS_GROUPS_REACHED**

  The domain exceeds the maximum allowed number of consistency groups.
- **DOMAIN_MAX_VOLUMES_REACHED**

  The domain exceeds the maximum allowed number of volumes.
- **OPERATION_DENIED_OBJECT_MANAGED**

  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.
- **SOURCE_AND_DESTINATION_DOMAINS_MUST_BE_DIFFERENT**

  Source and destination domains must be different.
- **NO_SPACE**

  The system does not have enough free space for the requested Storage Pool size

Setting the domain attribute

Use the **domain_manage** command to set or clear the Managed attribute of a domain.

```
domain_manage domain=DomainName managed=<yes|no>
```

Parameters

Name	Type	Description	Mandatory
domain	Object name	The domain name.	Y
managed	Boolean	The Managed attribute. If set to Yes, the domain will be marked as managed.	Y

Example:

```
domain_mmanage domain=d1 managed=yes
```

Output:

```
Command 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_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
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

User Category	Permission
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
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

Example:

```
appadmin_capabilities_set value=basic
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission
Storage administrator	Disallowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Disallowed
Technicians	Disallowed

Getting the application administrator's scope of commands

Use the **appadmin_capabilities_get** command to display the state of the Application Administrator's capabilities.

```
appadmin_capabilities_get
```

Example:

```
appadmin_capabilities_get
```

Output:

```
BASIC
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Allowed
Read-only users	Allowed
Technicians	Allowed

Chapter 19. 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.

```
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)

© Copyright IBM Corp. 2017

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

User Category	Permission
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Allowed

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            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

Field ID	Field output	Default position
configured_rate	Configured Rate (GBaud)	N/A
current_rate	Current Rate (GBaud)	8
port_state	Port State	9
link_type	Link Type	10
error_count	Error Count	11
active_firmware	Active Firmware	12
credit	Credit	N/A
hba_vendor	HBA Vendor	N/A
is_enabled	Enabled	N/A
module	Module	N/A
serial	Serial	N/A
temperature	Temperature	N/A
part_number	Part Number	N/A
original_serial	Original Serial	N/A
model	Model	N/A
original_model	Original Model	N/A
requires_service	Requires Service	N/A
service_reason	Service Reason	N/A
start_statstime	Timestamp of Link Statistics	N/A
link_failure	Link Failure	N/A
loss_of_sync	Loss of Sync	N/A
loss_of_signal	Loss of Signal	N/A
primit_seq_prot_error	Primitive Sequence Protocol Error	N/A
invalid_tx_word	Invalid Transmission Word	N/A
invalid_crc	Invalid CRC	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing FC port tests

Use the **fc_port_test_list** command to list the status of the system's FC port tests.

```
fc_port_test_list [ fc_port=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
fc_port	Lists only a specific port.	N	All ports in all modules.

This command lists all or some FC port tests running on the system. When no parameters are specified, all tests are listed. If a port is specified, a single test is listed.

Example:

```
fc_port_test_list
```

Output:

```
Component ID     Status       Start Time   Duration (Sec)   Transfer Length   CRC
---------------- ------------ ------------ ---------------- ----------------- -----
1:FC_Port:14:1   not active   NA           0                0                 0
1:FC_Port:14:2   not active   NA           0                0                 0
1:FC_Port:14:3   not active   NA           0                0                 0
1:FC_Port:14:4   not active   NA           0                0                 0
1:FC_Port:3:1    not active   NA           0                0                 0
1:FC_Port:3:2    not active   NA           0                0                 0
1:FC_Port:3:3    not active   NA           0                0                 0
1:FC_Port:3:4    not active   NA           0                0                 0
1:FC_Port:5:1    not active   NA           0                0                 0
1:FC_Port:5:2    not active   NA           0                0                 0
1:FC_Port:5:3    not active   NA           0                0                 0
1:FC_Port:5:4    not active   NA           0                0                 0
1:FC_Port:6:1    not active   NA           0                0                 0
1:FC_Port:6:2    not active   NA           0                0                 0
1:FC_Port:6:3    not active   NA           0                0                 0
1:FC_Port:6:4    not active   NA           0                0                 0

Cont.:

Frame Length Error   Disparity   Loopback Mode   Data Pattern   Data Size   Test Count
-------------------- ----------- --------------- -------------- ----------- ------------
0                    0           internal        0000           0           0
0                    0           internal        0000           0           0
0                    0           internal        0000           0           0
0                    0           internal        0000           0           0
0                    0           internal        0000           0           0
0                    0           internal        0000           0           0
0                    0           internal        0000           0           0
0                    0           internal        0000           0           0
0                    0           internal        0000           0           0
0                    0           internal        0000           0           0
0                    0           internal        0000           0           0
0                    0           internal        0000           0           0
0                    0           internal        0000           0           0
0                    0           internal        0000           0           0
0                    0           internal        0000           0           0
0                    0           internal        0000           0           0

Cont.:

Abort On Error
----------------
no
no
no
no
no
no
no
no
no
no
no
no
no
no
no
no
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
start_time	Start Time	3
duration	Duration (Sec)	4
transfer_length	Transfer Length	5
crc	CRC	6

Field ID	Field output	Default position
frame_length_error	Frame Length Error	7
disparity	Disparity	8
loopback_mode	Loopback Mode	9
pattern	Data Pattern	10
data_size	Data Size	11
test_count	Test Count	12
increment	Test Increment	N/A
abort_on_error	Abort On Error	13

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Starting an FC port test

Use **fc_port_test_start** to run a test of the FC port.

```
fc_port_test_start fc_port=ComponentId [ loopback=<internal|external> ] [ pattern=PATTERN ]
 [ data_size=DATA_SIZE ] [ frames=FRAMES ] [ increment=INCREMENT ]
 [ abort_on_error=<no|yes> ]
```

Parameters

Name	Type	Description	Mandatory	Default
fc_port	N/A	Port identifier.	Y	N/A
loopback	Enumeration	Loopback mode to use.	N	internal
pattern	String	Data pattern.	N	1414
data_size	Positive integer	Data size.	N	2048
frames	Positive integer	Number of frames.	N	10000
increment	Positive integer	Test increment.	N	1
abort_on_error	Boolean	Abort the test in case of an error.	N	no

Example:

```
xcli.py fc_port_test_start fc_port=1:FC_Port:1:1
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Allowed

Return codes

- **COMPONENT_IS_NOT_AN_FC_PORT**
  Component must specify an FC port
- **FC_PORT_DOES_NOT_EXIST**
  Specified FC port does not exist
- **COMPONENT_DOES_NOT_EXIST**
  Component does not exist
- **TEST_NOT_ALLOWED_IN_CURRENT_STATUS**
  Component cannot be tested in its current status
- **FC_PORT_TEST_IN_PROGRESS**
  FC port test is already in progress
- **INVALID_FC_PORT_TEST_DATA_PATTERN**
  Data pattern for FC port test is invalid

Aborting an FC port test

Use **fc_port_test_abort** to abort a currently running FC port test.

```
fc_port_test_abort fc_port=ComponentId
```

Parameters

Name	Description	Mandatory
fc_port	Port identifier.	Y

Example:

```
xcli.py fc_port_test_abort fc_port=1:FC_Port:1:1
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed

User Category	Permission
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Allowed

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
- **COMPONENT_DOES_NOT_EXIST**
  Component does not exist
- **NO_FC_PORT_TEST_IN_PROGRESS**
  FC port is not running test

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

Displaying the InfiniBand connectivity status

Use the **pod_ib_connectivity_list** command to display the pod's InfiniBand link connectivity status.

```
pod_ib_connectivity_list [ module=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
module	Defines whether to display the InfiniBand link connectivity status for a specific module or for all modules.	N	All modules

Example:

```
pod_ib_connectivity_list
```

Output:

```
Module     module-1       module-2       module-3    Flash-Canister-1   Flash-Canister-2
---------- -----------    -----------    ----------- ------------------ ------------------
module-1   connected      disconnected   connected   connected          connected
module-2   disconnected   connected      connected   connected          connected
module-3   connected      connected      connected   connected          connected
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Chapter 20. Flash enclosure maintenance commands

This section describes the command-line interface (CLI) for maintaining flash enclosures.

Listing Flash enclosure status

Use the **flash_enclosure_list** command to list special flash enclosure statuses.

```
flash_enclosure_list [ flash_enclosure=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
flash_enclosure	Flash enclosure for which special statuses are to be listed.	N	All Flash Enclosures.

This command lists the status of each Flash enclosure, including:

- Component generic status
- Online canister
- Total number of canisters

Example:

```
flash_enclosure_list
```

Output:

```
Component ID          Status   Currently Functioning   Control Path Status
--------------------- -------- ----------------------- --------------------
1:Flash_Enclosure:4   OK       yes                     OK

Cont.:

Cluster IP      Redundancy State   FW level        Has Spare   Array Rebuild Percentage
--------------- ------------------ --------------- ----------- ------------------------
14.10.204.239   online             1.4.1.0-195.0   yes         None


Cont.:

Encryption State
----------------
Enrolled
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
control_path_status	Control Path Status	4
cluster_ip	Cluster IP	5

© Copyright IBM Corp. 2017

Field ID	Field output	Default position
redundancy_state	Redundancy State	6
fw_level	FW level	7
has_spare	Has Spare	8
array_rebuild_percentage	Array Rebuild Percentage	9
encryption_state	Encryption State	10
array_status	Array Status	N/A
fru_part_number	FRU Part Number	N/A
fru_identity	FRU Identity	N/A
temperature_state	Temperature State	N/A
required_service	Requires Service	N/A
service_reason	Service Reason	N/A
enabled	Enabled	N/A
cluster_id	Cluster ID	N/A
serial_number	Serial Number	N/A
encrypted	Encrypted	N/A
key_needed	Key Needed	N/A
base_guid	Base GUID	N/A
charging	Charging	N/A
flash_status	Flash Status	N/A
fw_upgrade_status	FW Upgrade Status	N/A
fw_upgrade_progress	FW Upgrade Progress	N/A
fw_upgrade_committed	FW Upgrade Committed	N/A
target_fw_version	Target FW. Version	N/A
fw_file_name	FW File Name	N/A
utility_file_name	Utility File Name	N/A
cr_key_last_modified_time	CR Key Last Modified Time	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Allowed
Read-only users	Allowed
Technicians	Allowed

Listing a Flash canister status

Use the **flash_canister_list** command to list special Flash canister statuses.

```
flash_canister_list [ canister=ComponentId | flash_enclosure=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
canister	Flash Canister for which special statuses are to be listed.	N	All Flash canisters.
flash_enclosure	Flash Enclosure for which special statuses are to be listed.	N	All Flash enclosures.

This command lists the statuses of the Flash canisters, including:

- Component generic status
- Canister ID
- Node ID
- Node name

Example:

```
flash_canister_list
```

Output:

```
Component ID           Status   Currently Functioning   Service IP      Raid Status
---------------------- -------- ---------------------- --------------- -------------
1:Flash_Canister:4:1   OK       yes                    14.10.204.205   OK
1:Flash_Canister:4:2   OK       yes                    14.10.204.238   OK

Cont.:

Control Path Status   Serial Connected   MgmtNode
--------------------- ------------------ ----------
OK                    1:Module:12        no
OK                    1:Module:13        yes
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
service_ip	Service IP	4
raid_status	Raid Status	5
control_path_status	Control Path Status	6
serial_connected	Serial Connected	7
active	MgmtNode	8
fru_part_number	FRU Part Number	N/A
fru_identity	FRU Identity	N/A
temperature_state	Temperature State	N/A
fw_level	FW Level	N/A
mac_addresses	MAC Addresses	N/A
required_service	Requires Service	N/A
service_reason	Service Reason	N/A
canister_degraded	Canister Degraded	N/A
canister_missing	Canister Missing	N/A
status_led	Status Led	N/A

Field ID	Field output	Default position
check_log_led	Check Log Led	N/A
identify_led	Identify Led	N/A
controller_fault_led	Controller Fault Led	N/A
fault_led	Fault Led	N/A
dump_led	Dump Led	N/A
canister_mode	Canister Mode	N/A
service_mode	Service Mode	N/A
miswired	Miswired	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Allowed
Read-only users	Allowed
Technicians	Allowed

Listing a Flash card status

Use the **flash_card_list** command to list special Flash card statuses.

```
flash_card_list [ flash_card=ComponentId | flash_enclosure=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
flash_card	Flash card for which special statuses are to be listed.	N	All Flash cards.
flash_enclosure	Flash enclosure for which special statuses are to be listed.	N	All Flash enclosures.

This command lists the statuses of the TMS, including:

- Component generic status
- Slot ID
- Capacity
- Health State
- Flash type

Example:

```
flash_card_list flash_card=1:Flash_Card:4:5
```

Output:

```
Component ID       Status   Currently Functioning   Slot ID   Capacity(GB)   Health State
------------------ -------- ---------------------- --------- -------------- --------------
1:Flash_Card:4:5   OK       yes                    5         5717           good

Cont.:

Usage    Missing
-------- ---------
member   no
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
slot_id	Slot ID	4
capacity	Capacity(GB)	5
health_state	Health State	6
drive_use	Usage	7
missing	Missing	8
fru_part_number	FRU Part Number	N/A
fru_identity	FRU Identity	N/A
temperature_state	Temperature State	N/A
fw_level	FW Level	N/A
required_service	Requires Service	N/A
service_reason	Service Reason	N/A
flash_type	Type	N/A
fault_led	Fault LED	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Allowed
Read-only users	Allowed
Technicians	Allowed

Retrieving the list of Flash fans

Use the **flash_fan_list** command to retrieve the list of Flash fans.

```
flash_fan_list  [ flash_fan=ComponentId | flash_enclosure=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
flash_fan	Flash fan component ID	N	all
flash_enclosure	Flash enclosure component ID	N	all

Example:

```
flash_fan_list
```

Output:

```
Component ID      Status
----------------- --------
1:Flash_Fan:2:1   OK
1:Flash_Fan:2:2   OK
1:Flash_Fan:2:3   OK
1:Flash_Fan:2:4   OK
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
fru_part_number	FRU Part Number	N/A
fru_identity	FRU Identity	N/A
temperature_state	Temperature State	N/A
fw_level	FW Level	N/A
required_service	Requires Service	N/A
service_reason	Service Reason	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Retrieving the list of Flash IB adapters

Use the **flash_ib_adapter_list** command to retrieve the list of Flash IB adapters.

```
flash_ib_adapter_list  [ flash_ib_adapter=ComponentId | flash_enclosure=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
flash_ib_adapter	Flash IB adapter component ID	N	all
flash_enclosure	Flash Enclosure component ID	N	all

Example:

```
flash_ib_adapter_list
```

Output:

```
Component ID             Status
------------------------ --------
1:Flash_IB_Adapter:2:1   OK
1:Flash_IB_Adapter:2:2   OK
1:Flash_IB_Adapter:2:3   OK
1:Flash_IB_Adapter:2:4   OK
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
canister_id	Canister_name	4
fru_part_number	FRU Part Number	N/A
fru_identity	FRU Identity	N/A
temperature_state	Temperature State	N/A
required_service	Requires Service	N/A
service_reason	Service Reason	N/A
fw_level	FW Level	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Retrieving the Flash control connectivity list

Use the **flash_control_connectivity_list** command to retrieve the Flash control connectivity list.

```
flash_control_connectivity_list  [ canister=ComponentId | flash_enclosure=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
canister	Canister component ID	N	all
flash_enclosure	Flash enclosure Component ID	N	all

The connectivity for Flash control is provided via Ethernet.

Example:

```
flash_control_connectivity_list
```

Output:

```
Component ID           Path1        IPAddr1       P1Status   Path2        IPAddr2
---------------------- ------------ ------------- ---------- ------------ -------------
1:Flash_Canister:1:1   1:Module:1   14.10.204.1   yes        1:Module:2   14.10.204.33
1:Flash_Canister:1:2   1:Module:1   14.10.204.2   yes        1:Module:2   14.10.204.34

P2Status   Path3        IPAddr3        P3Status   Path4        IPAddr4         P4Status
---------- ------------ -------------- ---------- ------------ --------------- ----------
yes        1:Module:3   14.10.204.65   yes        1:Module:4   14.10.204.97    yes
yes        1:Module:3   14.10.204.66   yes        1:Module:4   14.10.204.98    yes
```

Field ID	Field output	Default position
component_id	Component ID	1
virtual_ips.0.path	Path1	2
virtual_ips.0.pathAddr	PathAddr1	3
virtual_ips.0.status	P1Status	4
virtual_ips.0.state	P1State	N/A
virtual_ips.1.path	Path2	5
virtual_ips.1.pathAddr	PathAddr2	6
virtual_ips.1.status	P2Status	7
virtual_ips.1.state	P2State	N/A
virtual_ips.2.path	Path3	8
virtual_ips.2.pathAddr	PathAddr3	9
virtual_ips.2.status	P3Status	10
virtual_ips.2.state	P3State	N/A
virtual_ips.3.path	Path4	11
virtual_ips.3.pathAddr	PathAddr4	12
virtual_ips.3.status	P4Status	13
virtual_ips.3.state	P4State	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Retrieving the list of Flash PSUs

Use the **flash_psu_list** command to retrieve the list of Flash PSUs.

```
flash_psu_list  [ flash_psu=ComponentId | flash_enclosure=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
flash_psu	Flash PSU component ID	N	all
flash_enclosure	Flash Enclosure component ID	N	all

Example:

```
flash_psu_list
```

Output:

```
Component ID      Status
----------------- --------
1:Flash_PSU:2:1   OK
1:Flash_PSU:2:2   OK
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
input_failed	Input Failed	4
output_failed	Output Failed	5
fru_part_number	FRU Part Number	N/A
fru_identity	FRU Identity	N/A
temperature_state	Temperature State	N/A
fw_level	FW Level	N/A
required_service	Requires Service	N/A
service_reason	Service Reason	N/A
fan_failed	Fan Failed	N/A
fault_led	Fault LED	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Retrieving the list of Flash PIBs

Use the **flash_pib_list** command to retrieve the list of Flash PIBs.

```
flash_pib_list [ flash_pib=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
flash_pib	Flash PIB component ID	N	all

Example:

```
flash_pib_list
```

Output:

```
Component ID      Status
----------------- --------
1:Flash_PIB:2:1   OK
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
fru_part_number	FRU Part Number	N/A
fru_identity	FRU Identity	N/A
temperature_state	Temperature State	N/A
fw_level	FW Level	N/A
required_service	Requires Service	N/A
service_reason	Service Reason	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Retrieving the list of Flash LED cards

Use the **flash_led_card_list** command to retrieve the list of Flash LED cards.

```
flash_led_card_list [ flash_led_card=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
flash_led_card	Flash LED card component ID	N	all

Example:

```
flash_led_card_list
```

Output:

```
Component ID      Status
----------------- --------
1:Flash_LED_Card:2:1   OK
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
fru_part_number	FRU Part Number	N/A
fru_identity	FRU Identity	N/A
temperature_state	Temperature State	N/A
fw_level	FW Level	N/A
required_service	Requires Service	N/A
service_reason	Service Reason	N/A
power_led	Power Led	N/A
fault_led	Fault Led	N/A
check_log_led	Check Log Led	N/A
identify_led	Identify Led	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing Flash BBU components

Use the **flash_bbu_list** command to list Flash BBU components.

```
flash_bbu_list [ flash_bbu=ComponentId | flash_enclosure=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
flash_bbu	Flash BBU component ID	N	all
flash_enclosure	Flash Enclosure component ID	N	all

Example:

```
flash_bbu_list
```

Output:

```
Component ID      Status   Charging Status   Percent Charged   Recondition Needed
----------------- -------- ----------------- ----------------- --------------------
1:Flash_BBU:2:1   OK       idle              94                no
1:Flash_BBU:2:2   OK       idle              89                no
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
percent_charged	Percent Charged	4
recondition_needed	Recondition Needed	5
charging_status	Charging Status	6
last_recondition_time	Last Recondition Time	7
temperature_state	Temperature State	N/A
fw_level	FW Level	N/A
required_service	Requires Service	N/A
service_reason	Service Reason	N/A
fru_part_number	FRU Part Number	N/A
fru_identity	FRU Identity	N/A
fault_led	Fault LED	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Retrieving the serial number of a Flash enclosure 1S

Use the **flash_enclosure_1s_get** command to retrieve the serial number of a Flash enclosure 1S.

```
flash_enclosure_1s_get flash_enclosure=ComponentId
```

Parameters

Name	Description	Mandatory
flash_enclosure	Flash enclosure component ID	Y

Example:

```
flash_enclosure_ls_get
```

Field ID	Field output	Default position
full_serial	Serial	1
mtm	Mtm	N/A
serial	Serial	N/A

Example:

```
flash_enclosure_ls_get
```

Output:

```
 Serial
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Chapter 21. Hardware maintenance commands

This section describes the command-line interface (CLI) for maintaining hardware components.

Monitoring the redistribution process

Use the **monitor_redist** command to monitor the status of the redistribution process.

```
monitor_redist
```

This command outputs the current redistribution process running on a module. The command does not yield information about a Flash enclosure.

The command may inform you that no such process exists. If such a process exists, the following information is shown:

- Type (adding new capacity, replacing failed component, phase-out, redistribution after failure)
- Initial capacity to copy
- Time started
- Capacity remaining to copy
- Time elapsed
- Percent completed
- Estimated time to completion

Field ID	Field output	Default position
type	Type	1
partitions_total	Initial Partitions	2
partitions_left	Partitions Remaining	3
percent_done	% Done	4
time_started	Time Started	5
estimated_time_to_finish	Estimated Time to Finish	6
time_elapsed	Time Elapsed	7

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

© Copyright IBM Corp. 2017

Displaying the system's average power consumption

Use the **system_average_power_consumption** command to display the system's average power consumption.

```
system_average_power_consumption
```

Field ID	Field output	Default position
value	Value	1
timestamp	Timestamp	2

Example:

```
system_average_power_consumption
```

Output:

```
Value       Timestamp
----------- -----------
1553        31-Aug-2016 16:25:01
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Getting the values for calculating the system's average power consumption

Use the **system_average_power_prepare** command to fetch all the values needed to calculate the system's average power consumption.

```
system_average_power_prepare
```

Example:

```
system_average_power_prepare
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Allowed

Return codes

- **HOT_UPGRADE_IS_IN_PROGRESS**

  Command is not allowed during hot upgrade

  Troubleshooting: Wait for hot upgrade to complete and try again
- **ALREADY_GETTING_FLASH_CANISTER_SNAP**

  Already getting flash canister snap
- **FAILED_GETTING_FLASH_CANISTER_SNAP**

  Failed getting flash canister snap
- **COMPONENT_DOES_NOT_EXIST**

  Component does not exist
- **COMMAND_IS_NOT_VALID_DURING_FLASH_FW_UPDATE**

  The requested command cannot be invoked while flash system update is running
- **SYSTEM_POWER_PREPARE_ALREADY_IN_PROGRESS**

  There is already a system power prepare command in progress

Displaying the system's average temperature

Use the **system_average_temperature** command to display the system's average temperature.

```
system_average_temperature
```

Field ID	Field output	Default position
value	Value	1
timestamp	Timestamp	2

Example:

```
system_average_temperature
```

Output:

```
Value       Timestamp
----------- --------------------
17          31-Aug-2016 16:21:51
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Enabling XIV Support access

Use the **xiv_support_enable** command to enable XIV Support access for a specific period of time limiting access from the specific address.

```
xiv_support_enable [ start=TimeStamp ]
 < finish=TimeStamp | timeout=Timeout > from=<IPaddress1[,IPaddress2]...> comment=Comment
```

Parameters

Name	Type	Description	Mandatory	Default
start	N/A	Start time for allowing XIV Support access.	N	Immediately.
finish	N/A	End time for allowing XIV Support access.	N	N/A
timeout	N/A	Timeout for allowing XIV Support access in either hh:mm format, or a number of minutes. The timeout cannot exceed 23 hours and 59 minutes. The word `unlimited` denotes unexpired timeout.	N	N/A
from	N/A	The source address to which XIV Support access is limited. It may be either IPv4 or IPv6 address, or `any`, or `technician` denoting laptop port.	Y	N/A
comment	String	Reason why XIV Support access is enabled.	Y	N/A

This command enables XIV Support access for a specific period of time limiting access from the specific address.

Example:

```
xiv_support_enable finish=2012-2-3.16:30 from=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 a positive number and define a 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 a 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
```

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
start	Start	2
finish	Finish	3
comment	Comment	4
from.0	From 0	5
from.1	From 1	6
from.2	From 2	7
from.3	From 3	8
from.4	From 4	9

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Allowed
Read-only users	Allowed
Technicians	Allowed

Listing system components

Use the **component_list** command to list system components and their status.

```
component_list [ component=ComponentId ] [ filter=<ALL|FAILED|NOTOK> ]
```

Parameters

Name	Type	Description	Mandatory	Default
component	N/A	Lists only this component.	N	All components.
filter	Enumeration	Filters the list to show only failed or only non-OK components.	N	ALL

The list can be filtered to show only a specific component, all failed components or all components in a non-OK state.

For status and configuration of specific component types, refer to the **_list** commands for specific components, such as: **module_list** or **switch_list**.

Example:

```
component_list
```

Output:

```
Component ID               Status         Currently Functioning
-------------------------- -------------- ----------------------
1:BBU:12:1                 Failed         no
1:Boot_Media:12:1          OK             yes
1:Boot_Media:12:2          OK             yes
1:CNA:12:1                 OK             yes
1:CNA:13:1                 OK             yes
1:CNA:8:1                  OK             yes
1:CPU:12:1                 OK             yes
1:CPU:12:2                 OK             yes
1:DIMM:12:1                OK             yes
1:DIMM:12:10               OK             yes
1:Data:12                  OK             yes
1:Data:13                  OK             yes
1:Data:8                   OK             yes
1:Data_Reduction:12        OK             yes
1:Data_Reduction:13        OK             yes
1:Data_Reduction:8         OK             yes
1:FC_Port:12:1             OK             yes
1:FC_Port:12:2             OK             yes
1:Fan:12:1                 OK             yes
1:Fan:12:2                 OK             yes
1:Flash_BBU:4:1            OK             yes
1:Flash_BBU:4:2            OK             yes
1:Flash_Canister:4:1       OK             yes
1:Flash_Canister:4:2       OK             yes
1:Flash_Card:4:1           Failed         no
1:Flash_Enclosure:4        OK             yes
1:Flash_Fan:4:1            OK             yes
1:Flash_Fan:4:2            OK             yes
1:Flash_Fan:4:3            OK             yes
1:Flash_Fan:4:4            OK             yes
1:Flash_IB_Adapter:4:1     OK             yes
1:Flash_IB_Adapter:4:2     OK             yes
1:Flash_IB_Adapter:4:3     OK             yes
1:Flash_IB_Adapter:4:4     OK             yes
1:Flash_LED_Card:4:1       OK             yes
1:Flash_PIB:4:1            OK             yes
1:Flash_PSU:4:1            OK             yes
1:Flash_PSU:4:2            OK             yes
1:IB_FlashSystem_Port:4:1  OK             yes
1:IB_FlashSystem_Port:4:3  OK             yes
1:IB_FlashSystem_Port:4:5  OK             yes
1:IB_FlashSystem_Port:4:7  OK             yes
1:IB_Module_Port:12:1      OK             yes
```

```
Cont:

1:IB_Module_Port:12:2       Failed         no
1:IB_Module_Port:13:1       OK             yes
1:IB_Module_Port:13:2       Failed         no
1:IB_Module_Port:8:1        OK             yes
1:IB_Module_Port:8:2        Failed         no
1:IB_Switch:1               OK             yes
1:IB_Switch:2               OK             yes
1:IB_Switch_BBU:1:1         Initializing   yes
1:IB_Switch_BBU:1:2         Initializing   yes
1:IB_Switch_BBU:2:1         Initializing   yes
1:IB_Switch_BBU:2:2         Initializing   yes
1:IB_Switch_Fan:1:1         Initializing   yes
1:IB_Switch_PSU:1:1         Initializing   yes
1:IB_Switch_PSU:1:2         Initializing   yes
1:IB_Switch_PSU:2:1         Initializing   yes
1:IB_Switch_PSU:2:2         Initializing   yes
1:IB_Switch_Port:1:12       OK             yes
1:IB_Switch_Port:1:13       OK             yes
1:IB_Switch_Port:1:20       OK             yes
1:Interface:12              OK             yes
1:Interface:13              OK             yes
1:Interface:8               OK             yes
1:Module:12                 OK             yes
1:Module:13                 OK             yes
1:Module:8                  OK             yes
1:NIC:12:1                  OK             yes
1:NIC:12:2                  OK             yes
1:PSU:12:1                  OK             yes
1:PSU:12:2                  OK             yes
1:PSU:13:1                  OK             yes
1:PSU:13:2                  OK             yes
1:PSU:8:1                   OK             yes
1:PSU:8:2                   OK             yes
1:Remote:12                 OK             yes
1:Remote:13                 OK             yes
1:Remote:8                  OK             yes
1:Vault_Device:12:1         OK             yes
1:Vault_Device:12:2         OK             yes
1:Vault_Device:13:1         OK             yes
1:Vault_Device:13:2         OK             yes
1:Vault_Device:8:1          OK             yes
1:Vault_Device:8:2          OK             yes
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
requires_service	Requires Service	N/A
service_reason	Service Reason	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing module configuration

Use the **module_list** command to list the configuration of all or specified modules.

```
module_list [ module=ModuleNumber ]
```

Parameters

Name	Description	Mandatory	Default
module	Lists the configuration of the specified module.	N	All modules

The following modules are available in FlashSystem A9000 and FlashSystem A9000R:

- BBU, also referred to or known as battery module
- Boot media, also referred to or known as hard disk drive (HDD)
- Compute module, also referred to or known as grid controller or module
- IB Switch, also referred to or known as InfiniBand switch
- Module, also referred to or known as grid controller or compute module.

This command lists the following information for each module:

- Generic component status
- Module type
- Number of boot media
- Number of vault devices
- Number of FC ports
- Number of Ethernet ports for iSCSI

Additional information is available through running `module_list -t all` :

- Serial
- Original serial
- Part number
- Original part number

Example:

```
module_list
```

Output:

```
Component ID   Status   Currently Functioning   Target Status   Type
-------------- -------- ----------------------- --------------- ------------------------
1:Module:12    OK       yes                                     g4.0_compute_enclosure
1:Module:13    OK       yes                                     g4.0_compute_enclosure
1:Module:8     OK       yes                                     g4.0_compute_enclosure

Cont.:

Boot Media Disks   Vault Devices   FC Ports   iSCSI Ports   Temperature
------------------ --------------- ---------- ------------- -------------
2                  2               4          2             22
2                  2               4          2             22
2                  2               4          2             22
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
target_status	Target Status	4
type	Type	5
boot_media_disks	Boot Media Disks	6
vault_devices	Vault Devices	7
fc_port_count	FC Ports	8
ethernet_port_count	iSCSI Ports	9
temperature	Temperature	10
avg_power	Avg Power	N/A
serial	Serial	N/A
original_serial	Original Serial	N/A
part_number	Part Number	N/A
original_part_number	Original Part Number	N/A
sas_version	SAS	N/A
infiniband_hca_version.0	InfiniBand HCA 1	N/A
infiniband_hca_version.1	InfiniBand HCA 2	N/A
cna_version.0	CNA 1	N/A
cna_version.1	CNA 2	N/A
compression_adapter_firmware.0	Compression Adapter 1	N/A
compression_adapter_firmware.1	Compression Adapter 2	N/A
requires_service	Requires Service	N/A
service_reason	Service Reason	N/A
memory_gb	Mem	N/A
module_11s_number	11S Number	N/A
megaraid_serial_number	MegaRAID Serial	N/A
megaraid_product_name	MegaRAID Product Name	N/A
megaraid_package_version	MegaRAID Package Version	N/A
megaraid_flash_components.0	MegaRAID Flash Component 1	N/A
megaraid_flash_components.1	MegaRAID Flash Component 2	N/A
megaraid_flash_components.2	MegaRAID Flash Component 3	N/A
megaraid_flash_components.3	MegaRAID Flash Component 4	N/A
megaraid_flash_components.4	MegaRAID Flash Component 5	N/A
megaraid_flash_components.5	MegaRAID Flash Component 6	N/A
megaraid_flash_components.6	MegaRAID Flash Component 7	N/A
megaraid_flash_components.7	MegaRAID Flash Component 8	N/A
imm_version	IMM Version	N/A
uefi_version	UEFI Version	N/A
dsa_version	DSA Version	N/A
mcu_version	MCU Version	N/A
board_serial	Board Serial	N/A
board_part_number	Board Part Number	N/A
board_mfg_date	Board MFG Date	N/A
dasd_board_serial	Backplane Serial	N/A
dasd_board_part_number	Backplane Part Number	N/A

Field ID	Field output	Default position
dasd_board_manufacturer	Backplane Manufacturer	N/A
dasd_board_mfg_date	Backplane MFG Date	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing the internal temperature of modules

Use the **module_temperature_list** command to list the modules' internal temperatures in the storage system.

```
module_temperature_list [ module=ModuleNumber ]
```

Parameters

Name	Description	Mandatory	Default
module	Limits the listing to a specific module.	N	All temperatures in all modules.

Example:

```
module_temperature_list -f all
```

Output:

```
xcli.py module_temperature_list -f all
Module        Ambient   PCH   RS1   RS2   MID1   MID2   RAID   DIMM AB VR   DIMM CD VR
------------- --------- ----- ----- ----- ------ ------ ------ ------------ ------------
1:Module:1    16        44    32    38    31     31     34     27           27
1:Module:11   16        46    35    40    34     33     35     26           28
1:Module:2    16        43    34    38    33     32     33     26           27
1:Module:4    16        43    34    37    33     32     34     25           26

DIMM EF VR   DIMM GH VR   CPU1   CPU2   InfiniBand HCA   PSU FR   Fibre Channel Port fc-0
------------ ------------ ------ ------ ---------------- -------- -------------------------
28           27           52     55     53               28       0
27           32           58     68     56               32       0
28           27           51     53     53               28       0
29           30           51     55     53               29       0

Fibre Channel Port fc-2    CNA=0   CNA-1   Boot Device 0   Boot Device 1   Vault Device 0
-------------------------  ------- ------- --------------- --------------- ----------------
0                          50      46      23              21              20
0                          51      52      27              25              21
0                          50      49      22              22              20
0                          48      52      24              23              20

Vault Device 1   BBU1   BBU2   CPU1 VR   CPU2 VR   Fibre Channel Port fc-1   Fibre Channel
                                                                            Port fc-3
---------------- ------ ------ --------- --------- ------------------------ -------------
21               18     18     35        36        0                        0
21               17     16     37        40        0                        0
20               18     17     35        37        0                        0
20               17     17     35        37        0                        0
```

Field ID	Field output	Default position
component_id	Module	1
sdr_temperatures.0	Ambient	2
sdr_temperatures.1	PCH	3
sdr_temperatures.2	RS1	4
sdr_temperatures.3	RS2	5
sdr_temperatures.4	MID1	6
sdr_temperatures.5	MID2	7
sdr_temperatures.6	RAID	8
sdr_temperatures.8	CPU1 VR	N/A
sdr_temperatures.9	CPU2 VR	N/A
sdr_temperatures.10	DIMM AB VR	9
sdr_temperatures.11	DIMM CD VR	10
sdr_temperatures.12	DIMM EF VR	11
sdr_temperatures.13	DIMM GH VR	12
sdr_temperatures.14	CPU1	13
sdr_temperatures.15	CPU2	14
ib_hca_temperature	InfiniBand HCA	15
sdr_temperatures.7	PSU FR	16
fc_adapter_temperature.0	Fibre Channel Port fc-0	17
fc_adapter_temperature.1	Fibre Channel Port fc-1	N/A
fc_adapter_temperature.2	Fibre Channel Port fc-2	18
fc_adapter_temperature.3	Fibre Channel Port fc-3	N/A
cna_temperature.0	CNA=0	19
cna_temperature.1	CNA-1	20
boot_device_temperature.0	Boot Device 0	21

Field ID	Field output	Default position
boot_device_temperature.1	Boot Device 1	22
vault_device_temperature.0	Vault Device 0	23
vault_device_temperature.1	Vault Device 1	24
bbu_temperature.0	BBU1	25
bbu_temperature.1	BBU2	26

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing boot media devices in the system

Use the **boot_media_list** command to list boot media devices in the storage system.

```
boot_media_list [ module=ModuleNumber | boot_media=BootMediaDevice ]
```

Parameters

Name	Description	Mandatory	Default
module	Limits the listing to the specific module.	N	Boot media devices in a specific module.
boot_media	Limits the listing to a specific boot media.	N	A specific boot media device.

Example:

```
boot_media_list -f all
```

Output:

```
Component ID        Status   Currently Functioning   Hardware Status   Vendor
------------------- -------- ---------------------- ----------------- --------
1:Boot_Media:11:1   OK       yes                    OK                IBM-ESXS
1:Boot_Media:11:2   OK       yes                    OK                IBM-ESXS
1:Boot_Media:1:1    OK       yes                    OK                IBM-ESXS
1:Boot_Media:1:2    OK       yes                    OK                IBM-ESXS
1:Boot_Media:2:1    OK       yes                    OK                IBM-ESXS
1:Boot_Media:2:2    OK       yes                    OK                IBM-ESXS
1:Boot_Media:4:1    OK       yes                    OK                IBM-ESXS
1:Boot_Media:4:2    OK       yes                    OK                IBM-ESXS

Model              Serial   FW     Temperature   Original Serial   Part #
------------------ -------- ------ ------------- ----------------- --------
HUC101860CS420 X   03V0E75K J5H2   26            03V0E75K
HUC101860CS420 X   03V0LPEK J5H2   25            03V0LPEK
HUC101860CS420 X   03V0X90H J5H2   21            03V0X90H
HUC101860CS420 X   03V0AHME J5H2   20            03V0AHME
HUC101860CS420 X   03V117XE J5H2   22            03V117XE
HUC101860CS420 X   03V0DAGK J5H2   21            03V0DAGK
HUC101860CS420 X   03V0DJ4K J5H2   23            03V0DJ4K
HUC101860CS420 X   03V0H38K J5H2   22            03V0H38K

Original Part #  Size (GB)   Requires Service   Service Reason
---------------- ---------- ------------------ ----------------
                 600 GB
                 600 GB
                 600 GB
                 600 GB
                 600 GB
                 600 GB
                 600 GB
                 600 GB

Rebuild Progress   Rebuild Time (sec.)
------------------ --------------------
N/A                N/A
N/A                N/A
N/A                N/A
N/A                N/A
N/A                N/A
N/A                N/A
N/A                N/A
N/A                N/A
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
hardware_status	Hardware Status	4
vendor	Vendor	5
model	Model	6
serial	Serial	7
fw_revision	FW	8
temperature	Temperature	N/A
original_serial	Original Serial	N/A
part_number	Part #	N/A
original_part_number	Original Part Number	N/A
size	Size (GB)	N/A
requires_service	Requires Service	N/A
service_reason	Service Reason	N/A
rebuild_progress	Rebuild Progress	N/A

Field ID	Field output	Default position
rebuild_time	Rebuild Time (sec.)	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

Viewing vault devices in the system

Use the **vault_device_list** command to view the status of special vault devices.

```
vault_device_list [ module=ModuleNumber  | vault_device=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
module	Limits the listing to a specific module.	N	All vault devices in all modules.
vault_device	Vault devices for which special statuses are to be listed.	N	All vault devices.

Example:

```
vault_device_list
```

Output:

```
Component ID          Status   Currently Functioning   Capacity   Target Status   Vendor
--------------------- -------- ---------------------- ---------- --------------- ----------
1:Vault_Device:10:1   OK       yes                    250GB                      LENOVO-X
1:Vault_Device:10:2   OK       yes                    250GB                      LENOVO-X
1:Vault_Device:11:1   OK       yes                    250GB                      LENOVO-X
1:Vault_Device:11:2   OK       yes                    250GB                      LENOVO-X
1:Vault_Device:7:1    OK       yes                    250GB                      LENOVO-X
1:Vault_Device:7:2    OK       yes                    250GB                      LENOVO-X
1:Vault_Device:9:1    OK       yes                    250GB                      LENOVO-X
1:Vault_Device:9:2    OK       yes                    250GB                      LENOVO-X

Cont.:

Model             Serial     Firmware   FRU       Temperature   Encryption State
----------------- ---------- ---------- --------- ------------- ----------------
HUSMR1625ASS20E   0PVGJTPA   P4C9       00NA685   22            Ready
HUSMR1625ASS20E   0PVGHN6A   P4C9       00NA685   22            Ready
HUSMR1625ASS20E   0PVJ2PEA   P4C9       00NA685   21            Ready
HUSMR1625ASS20E   0PVJ251A   P4C9       00NA685   21            Ready
HUSMR1625ASS20E   0PVJ99UA   P4C9       00NA685   20            Ready
HUSMR1625ASS20E   0PVJRB5A   P4C9       00NA685   21            Ready
HUSMR1625ASS20E   0PVJS39A   P4C9       00NA685   20            Ready
HUSMR1625ASS20E   0PVJ9RAA   P4C9       00NA685   21            Ready
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
capacity_in_bytes	Capacity	N/A
capacity	Capacity	4
target_status	Target Status	5
vendor	Vendor	6
original_vendor	Original Vendor	N/A
model	Model	7
original_model	Original Model	N/A
serial	Serial	8
original_serial	Original Serial	N/A
firmware	Firmware	9
original_firmware	Original Firmware	N/A
part_number	FRU	10
original_part_number	Original FRU	N/A
temperature	Temperature	11
encryption_state	Encryption State	12
requires_service	Requires Service	N/A
service_reason	Service Reason	N/A
revision	Revision	N/A
drive_pn	Drive P/N	N/A
original_drive_pn	Original Drive P/N	N/A
fru_pn	FRU P/N	N/A
original_fru_pn	Original FRU P/N	N/A
desc.bgd_scan	Background Scan	N/A
desc.disk_id	Device ID	N/A
desc.last_sample_serial	Last Sample Serial	N/A
desc.last_sample_time	Last Sample Time	N/A

Field ID	Field output	Default position
desc.power_is_on	Power On	N/A
desc.power_on_hours	Power On Hours	N/A
desc.power_on_minutes	Power On Minutes	N/A
desc.last_time_pom_was_mod	Last Time Power On Minutes Was Modified	N/A
desc.read_fail	Read Fail	N/A
desc.smart_code	SMART Code	N/A
desc.smart_fail	SMART Fail	N/A
desc.temperature_ status.reported_severity	Reported Temperature Severity	N/A
desc.temperature_ status.reported_temperature	Reported Temperature	N/A
desc.temperature_ status.temperature	Device Temperature	N/A
desc.sw_encryption_active	Software-Based Encryption Active	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Allowed
Read-only users	Allowed
Technicians	Allowed

Listing BBUs in the system

Use the **module_bbu_list** command to list the BBUs in the storage system modules.

```
module_bbu_list [ module=ModuleNumber | module_bbu=BbuNumber ]
```

Parameters

Name	Description	Mandatory	Default
module	Limits the listing to a specific module.	N	All BBUs in all modules.
module_bbu	Limits the listing to a specific BBU.	N	All BBUs.

Example:

```
module_bbu_list -f all
```

Output:

```
Component ID   Status   Curr Functioning   State   Hardware Status   Remaining Capacity
------------- -------- ------------------ ------- ----------------- --------------------
1:BBU:14:1     OK       yes                Full    OK                891
1:BBU:14:2     OK       yes                Full    OK                877
1:BBU:3:1      OK       yes                Full    OK                787
1:BBU:3:2      OK       yes                Full    OK                860
1:BBU:5:1      OK       yes                Full    OK                792
1:BBU:5:2      OK       yes                Full    OK                898
1:BBU:6:1      OK       yes                Full    OK                817
1:BBU:6:2      OK       yes                Full    OK                814

Cont.:

Full Charge Capacity   Charged %   Time to Empty   Time to Full   Charger State
---------------------- ----------- --------------- -------------- ---------------
891                    100         1600200         0              in progress
877                    100         0               0              in progress
787                    100         0               0              in progress
860                    100         0               0              in progress
792                    100         0               0              in progress
898                    100         0               0              in progress
817                    100         2944800         0              in progress
814                    100         0               0              in progress

Cont.:

Calibration State   Calibration Time
------------------- ------------------
Idle                0
Idle                0
Idle                0
Idle                0
Idle                0
Idle                0
Idle                0
Idle                0
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Curr Functioning	3
required_service	Requires Service	N/A
service_reason	Service Reason	N/A
bbu_state	State	4
bbu_status	Hardware Status	5
remaining_capacity	Remaining Capacity	6
full_charge_capacity	Full Charge Capacity	7
percent_charged	Charged %	8
time_to_empty	Time to Empty	9
time_to_full	Time to Full	10
charger_state	Charger State	11
calib_state	Calibration State	12
calib_requested	Calibration Needed	N/A
last_succ_calib_date	Successful Calibration Time	N/A
last_calib_date	Calibration Time	13
last_calib_result	Calibration Result	N/A
insertion_date	Inserted	N/A
manuf_date	Manufactured	N/A
fw	FW Version	N/A

Field ID	Field output	Default position
epow_cable_present	EPOW Cable Present	N/A
power_sense_cable_present	Power Sense Cable Present	N/A
epow_simulate	EPOW Simulate	N/A
epow_asserted	EPOW Asserted	N/A
cycle_count	Cycles	N/A
temperature_tenths_celsius	Temp /10C	N/A
charger_enabled	Charger Enabled	N/A
slow_charge_enabled	Slow Charge Enabled	N/A
discharge_enabled	Discharge Enabled	N/A
ps2_present	PS2 Present	N/A
charge_now	Nominal Available Capacity mAh	N/A
voltage_now	Voltage Now mV	N/A
current_now	Current Now mA	N/A
power_avg	Power Average mW	N/A
charge_full	Full Available Capacity mAh	N/A
charge_full_design	Design Charge	N/A
energy_now	Energy now uWh	N/A
at_rate	At Rate	N/A
at_rate_tte	At Rate Time to Empty	N/A
charge_now_sufficient	Charge Now Sufficient	N/A
endurance_start_monotonic_time	Endurance Start Monotonic Time	N/A
serial	Serial	N/A
original_serial	Original Serial	N/A
part_number	Part #	N/A
original_part_number	Original Part #	N/A
fru	FRU	N/A
runtime	Runtime	N/A
full_power_runtime	Full Power Runtime	N/A
half_power_runtime	Half Power Runtime	N/A
module_runtime	Module Runtime	N/A
state_of_health	Health	N/A
charge_voltage	Charge Voltage mV	N/A
charge_current	Charge Current mA	N/A
test_calib_en	Test/Calib. Enabled	N/A
fhd_enabled	FHD Enabled	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing PSUs in a module

Use the **module_psu_list** command to list PSUs in the specified module.

```
module_psu_list [ module=ModuleNumber  | psu=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
module	Limits the listing to a specific module.	N	All PSUs in all modules.
psu	Lists only a specific PSU.	N	A specific PSU.

Example:

```
xcli.py module_psu_list -f all
```

Output:

```
Component ID   Status   Currently Functioning   Location
-------------- -------- ---------------------- ----------------
1:PSU:1:1      Failed   no                     Power Supply 1
1:PSU:1:2      OK       yes                    Power Supply 2
1:PSU:2:1      OK       yes                    Power Supply 1
1:PSU:2:2      Failed   no                     Power Supply 2
1:PSU:4:1      OK       yes                    Power Supply 1
1:PSU:4:2      Failed   no                     Power Supply 2

Sensor status                              Serial number   Part number
------------------------------------------ --------------- -------------
Presence NOT detected                      N/A             N/A
Presence detected                          K115148J01L     94Y8143
Presence detected                          K115148B06E     94Y8143
Presence detected, Power Supply AC lost    K115148B01P     94Y8143
Presence detected                          K115148B0AP     94Y8143
Presence detected, Power Supply AC lost    K115148J01B     94Y8143

Requires Service   Service Reason
------------------ -------------------------
REPLACE            MODULE_PSU__NOT_DETECTED

COMPONENT_TEST     MODULE_PSU__BAD_POWER_INPUT

COMPONENT_TEST     MODULE_PSU__BAD_POWER_INPUT
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
location	Location	4
sensor_statuses	Sensor statuses	5
requires_service	Requires Service	N/A
service_reason	Service Reason	N/A
serial	Serial number	N/A
part_number	Part number	N/A
manufacturer	Manufacturer	N/A

Field ID	Field output	Default position
mfg_date	Manufacturing Date	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing compression adapters in the system

Use the **compression_adapter_list** command to list compression adapters in the system.

```
compression_adapter_list [ module=ModuleNumber  | compression_adapter=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
module	Limits the listing to a specific module.	N	All compression adapters in all modules.
compression_adapter	Lists only a specific compression adapter.	N	A specific compression adapter.

Compression adapters are used to increase the speed of I/O transfers to and from compressed volumes.

Example:

```
compression_adapter_list
```

Output:

```
Component ID               Status   Currently Functioning   Firmware   Hardware
-------------------------- -------- ----------------------- ---------- ----------
1:Compression_Adapter:6:1  OK       yes                     1.0.12     A0 SKU3
1:Compression_Adapter:6:2  OK       yes                     1.0.12     A0 SKU3

Driver
----------------
1.0.12 3d60d12
1.0.12 3d60d12
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
serial	Serial	4

Field ID	Field output	Default position
firmware_version	Firmware	5
hardware_version	Hardware	6
driver_version	Driver	7
type	Type	N/A
mmp_version	MMP	N/A
pci_address	PCI Address	N/A
quick_assist_api_cy_version	QuickAssist API CY	N/A
quick_assist_api_dc_version	QuickAssist API DC	N/A
threading_mode	Threading Mode	N/A
requires_service	Requires Service	N/A
service_reason	Service Reason	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing fans in a module

Use the **fan_list** command to list fans in the specified module.

```
fan_list [ module=ModuleNumber  | fan=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
module	Limits the listing to a specific module.	N	All fans in all modules.
fan	Lists only a specific fan.	N	A specific fan.

Example:

```
fan_list
```

Output:

```
Component ID   Status   Currently Functioning   Location   Speed
-------------- -------- ----------------------- ---------- -------
1:Fan:14:1     OK       yes                     1A         4012
1:Fan:14:10    OK       yes                     5B         3780
1:Fan:14:11    OK       yes                     6A         3953
1:Fan:14:12    OK       yes                     6B         3780
1:Fan:14:2     OK       yes                     1B         3717
1:Fan:14:3     OK       yes                     2A         4012
1:Fan:14:4     OK       yes                     2B         3780
1:Fan:14:5     OK       yes                     3A         4012
1:Fan:14:6     OK       yes                     3B         3780
1:Fan:14:7     OK       yes                     4A         4012
1:Fan:14:8     OK       yes                     4B         3780
1:Fan:14:9     OK       yes                     5A         3894
1:Fan:3:1      OK       yes                     1A         3894
1:Fan:3:10     OK       yes                     5B         3717
1:Fan:3:11     OK       yes                     6A         3953
1:Fan:3:12     OK       yes                     6B         3843
1:Fan:3:2      OK       yes                     1B         3717
1:Fan:3:3      OK       yes                     2A         4012
1:Fan:3:4      OK       yes                     2B         3780
1:Fan:3:5      OK       yes                     3A         3953
1:Fan:3:6      OK       yes                     3B         3780
1:Fan:3:7      OK       yes                     4A         3953
1:Fan:3:8      OK       yes                     4B         3717
1:Fan:3:9      OK       yes                     5A         3894
1:Fan:5:1      OK       yes                     1A         3953
1:Fan:5:10     OK       yes                     5B         3528
1:Fan:5:11     OK       yes                     6A         3953
1:Fan:5:12     OK       yes                     6B         3780
1:Fan:5:2      OK       yes                     1B         3780
1:Fan:5:3      OK       yes                     2A         3953
1:Fan:5:4      OK       yes                     2B         3780
1:Fan:5:5      OK       yes                     3A         3894
1:Fan:5:6      OK       yes                     3B         3780
1:Fan:5:7      OK       yes                     4A         3953
1:Fan:5:8      OK       yes                     4B         3780
1:Fan:5:9      OK       yes                     5A         3953
1:Fan:6:1      OK       yes                     1A         3953
1:Fan:6:10     OK       yes                     5B         3780
1:Fan:6:11     OK       yes                     6A         3894
1:Fan:6:12     OK       yes                     6B         3717
1:Fan:6:2      OK       yes                     1B         3780
1:Fan:6:3      OK       yes                     2A         3953
1:Fan:6:4      OK       yes                     2B         3654
1:Fan:6:5      OK       yes                     3A         3953
1:Fan:6:6      OK       yes                     3B         3843
1:Fan:6:7      OK       yes                     4A         3953
1:Fan:6:8      OK       yes                     4B         3780
1:Fan:6:9      OK       yes                     5A         4012
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
location_a	Location	N/A
rpm_a	Speed	4
min_rpm_a	Min Speed	N/A
max_rpm_a	Max Speed	N/A
location_b	Peer Location	N/A
rpm_b	Peer Speed	5
min_rpm_b	Peer Min Speed	N/A
max_rpm_b	Peer Max Speed	N/A
requires_service	Requires Service	N/A

Field ID	Field output	Default position
service_reason	Service Reason	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing NICs in the system

Use the **nic_list** command to list the NICs in the storage system.

```
nic_list [ module=ModuleNumber  | nic=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
module	Limits the listing to a specific module.	N	All NICs in all modules.
nic	Lists only a specific NIC.	N	A specific NIC.

Example:

```
nic_list -f all
```

Output:

```
root@nextra-6013742-module-9:~# xcli.py nic_list -f all
root@nextra-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 DIMMs in the modules

Use the **dimm_list** command to list the DIMMs in the modules.

```
dimm_list [ module=ModuleNumber  | dimm=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
module	Limits the listing to a specific module.	N	All DIMMs in all modules.
dimm	Lists only a specific DIMM.	N	A specific DIMM.

The memory modules (DIMMs) run the microcode and the data cache in the grid controller.

Example:

```
dimm_list -f all
```

Output:

```
Component ID   Status   Currently Functioning   Hardware Status   DIMM Id   CPU
-------------- -------- ---------------------- ----------------- --------- -----
1:DIMM:7:1     OK       yes                     OK                1         1
1:DIMM:7:10    OK       yes                     OK                10        1
1:DIMM:7:11    OK       yes                     OK                11        1
1:DIMM:7:12    OK       yes                     OK                12        1
1:DIMM:7:13    OK       yes                     OK                13        2
1:DIMM:7:14    OK       yes                     OK                14        2
1:DIMM:7:15    OK       yes                     OK                15        2
1:DIMM:7:16    OK       yes                     OK                16        2
1:DIMM:7:17    OK       yes                     OK                17        2
1:DIMM:7:18    OK       yes                     OK                18        2
1:DIMM:7:19    OK       yes                     OK                19        2
1:DIMM:7:2     OK       yes                     OK                2         1
1:DIMM:7:20    OK       yes                     OK                20        2
1:DIMM:7:21    OK       yes                     OK                21        2
1:DIMM:7:22    OK       yes                     OK                22        2
1:DIMM:7:23    OK       yes                     OK                23        2
1:DIMM:7:24    OK       yes                     OK                24        2
1:DIMM:7:3     OK       yes                     OK                3         1
1:DIMM:7:4     OK       yes                     OK                4         1
1:DIMM:7:5     OK       yes                     OK                5         1
1:DIMM:7:6     OK       yes                     OK                6         1
1:DIMM:7:7     OK       yes                     OK                7         1
1:DIMM:7:8     OK       yes                     OK                8         1
1:DIMM:7:9     OK       yes                     OK                9         1
1:DIMM:9:1     OK       yes                     OK                1         1
1:DIMM:9:10    OK       yes                     OK                10        1
1:DIMM:9:11    OK       yes                     OK                11        1
1:DIMM:9:12    OK       yes                     OK                12        1
1:DIMM:9:13    OK       yes                     OK                13        2
1:DIMM:9:14    OK       yes                     OK                14        2
1:DIMM:9:15    OK       yes                     OK                15        2
1:DIMM:9:16    OK       yes                     OK                16        2
1:DIMM:9:17    OK       yes                     OK                17        2
1:DIMM:9:18    OK       yes                     OK                18        2
1:DIMM:9:19    OK       yes                     OK                19        2
1:DIMM:9:2     OK       yes                     OK                2         1
1:DIMM:9:20    OK       yes                     OK                20        2
1:DIMM:9:21    OK       yes                     OK                21        2
1:DIMM:9:22    OK       yes                     OK                22        2
1:DIMM:9:23    OK       yes                     OK                23        2
1:DIMM:9:24    OK       yes                     OK                24        2
1:DIMM:9:3     OK       yes                     OK                3         1
1:DIMM:9:4     OK       yes                     OK                4         1
1:DIMM:9:5     OK       yes                     OK                5         1
1:DIMM:9:6     OK       yes                     OK                6         1
1:DIMM:9:7     OK       yes                     OK                7         1
1:DIMM:9:8     OK       yes                     OK                8         1
1:DIMM:9:9     OK       yes                     OK                9         1
```

```
Cont.:

Size(Mb)   Speed(MHz)   Configured Clock Speed(MHz)   Manufacturer   Serial
---------- ------------ ----------------------------- -------------- ----------
16384      2133         2133                          Hynix          505F63F9
16384      2133         2133                          Samsung        39542977
16384      2133         2133                          Samsung        3953EA6C
16384      2133         2133                          Samsung        3953F39A
16384      2133         2133                          Samsung        3953F240
16384      2133         2133                          Samsung        39542562
16384      2133         2133                          Samsung        3954300B
16384      2133         2133                          Samsung        39546472
16384      2133         2133                          Samsung        39540BB2
16384      2133         2133                          Samsung        3953FB59
16384      2133         2133                          Samsung        3954074A
16384      2133         2133                          Samsung        3953F241
16384      2133         2133                          Samsung        395404E0
16384      2133         2133                          Samsung        395425D8
16384      2133         2133                          Samsung        39542BF2
16384      2133         2133                          Samsung        395426EF
16384      2133         2133                          Samsung        3953EB61
16384      2133         2133                          Samsung        39542AD0
16384      2133         2133                          Samsung        39542973
16384      2133         2133                          Samsung        39542ACF
16384      2133         2133                          Samsung        3953E982
16384      2133         2133                          Samsung        39542568
16384      2133         2133                          Samsung        3953EA4A
16384      2133         2133                          Samsung        3953E993
16384      2133         2133                          Hynix          804AC8C2
16384      2133         2133                          Hynix          707387FA
16384      2133         2133                          Hynix          7073895C
16384      2133         2133                          Hynix          3077315A
16384      2133         2133                          Hynix          90655EDF
16384      2133         2133                          Hynix          7073885E
16384      2133         2133                          Hynix          30772789
16384      2133         2133                          Hynix          707388BB
16384      2133         2133                          Hynix          90655FC6
16384      2133         2133                          Hynix          70738960
16384      2133         2133                          Hynix          70738871
16384      2133         2133                          Hynix          7073881A
16384      2133         2133                          Hynix          307C97D6
16384      2133         2133                          Hynix          7073880D
16384      2133         2133                          Hynix          70738819
16384      2133         2133                          Hynix          30773136
16384      2133         2133                          Hynix          30772FF8
16384      2133         2133                          Hynix          70738850
16384      2133         2133                          Hynix          3079B2CD
16384      2133         2133                          Hynix          307A04D2
16384      2133         2133                          Hynix          70738951
16384      2133         2133                          Hynix          70738939
16384      2133         2133                          Hynix          70738A0B
16384      2133         2133                          Hynix          70738823
```

```
Cont.:

Original Serial    Part #             Original Part Number
-----------------  ------------------ ----------------------
505F63F9           HMA42GR7MFR4N-TF   HMA42GR7MFR4N-TF
39542977           M393A2G40DB0-CPB   M393A2G40DB0-CPB
3953EA6C           M393A2G40DB0-CPB   M393A2G40DB0-CPB
3953F39A           M393A2G40DB0-CPB   M393A2G40DB0-CPB
3953F240           M393A2G40DB0-CPB   M393A2G40DB0-CPB
39542562           M393A2G40DB0-CPB   M393A2G40DB0-CPB
3954300B           M393A2G40DB0-CPB   M393A2G40DB0-CPB
39546472           M393A2G40DB0-CPB   M393A2G40DB0-CPB
39540BB2           M393A2G40DB0-CPB   M393A2G40DB0-CPB
3953FB59           M393A2G40DB0-CPB   M393A2G40DB0-CPB
3954074A           M393A2G40DB0-CPB   M393A2G40DB0-CPB
3953F241           M393A2G40DB0-CPB   M393A2G40DB0-CPB
395404E0           M393A2G40DB0-CPB   M393A2G40DB0-CPB
395425D8           M393A2G40DB0-CPB   M393A2G40DB0-CPB
39542BF2           M393A2G40DB0-CPB   M393A2G40DB0-CPB
395426EF           M393A2G40DB0-CPB   M393A2G40DB0-CPB
3953EB61           M393A2G40DB0-CPB   M393A2G40DB0-CPB
39542AD0           M393A2G40DB0-CPB   M393A2G40DB0-CPB
39542973           M393A2G40DB0-CPB   M393A2G40DB0-CPB
39542ACF           M393A2G40DB0-CPB   M393A2G40DB0-CPB
3953E982           M393A2G40DB0-CPB   M393A2G40DB0-CPB
39542568           M393A2G40DB0-CPB   M393A2G40DB0-CPB
3953EA4A           M393A2G40DB0-CPB   M393A2G40DB0-CPB
3953E993           M393A2G40DB0-CPB   M393A2G40DB0-CPB
804AC8C2           HMA42GR7MFR4N-TF   HMA42GR7MFR4N-TF
707387FA           HMA42GR7MFR4N-TF   HMA42GR7MFR4N-TF
7073895C           HMA42GR7MFR4N-TF   HMA42GR7MFR4N-TF
3077315A           HMA42GR7MFR4N-TF   HMA42GR7MFR4N-TF
90655EDF           HMA42GR7MFR4N-TF   HMA42GR7MFR4N-TF
7073885E           HMA42GR7MFR4N-TF   HMA42GR7MFR4N-TF
30772789           HMA42GR7MFR4N-TF   HMA42GR7MFR4N-TF
707388BB           HMA42GR7MFR4N-TF   HMA42GR7MFR4N-TF
90655FC6           HMA42GR7MFR4N-TF   HMA42GR7MFR4N-TF
70738960           HMA42GR7MFR4N-TF   HMA42GR7MFR4N-TF
70738871           HMA42GR7MFR4N-TF   HMA42GR7MFR4N-TF
7073881A           HMA42GR7MFR4N-TF   HMA42GR7MFR4N-TF
307C97D6           HMA42GR7MFR4N-TF   HMA42GR7MFR4N-TF
7073880D           HMA42GR7MFR4N-TF   HMA42GR7MFR4N-TF
70738819           HMA42GR7MFR4N-TF   HMA42GR7MFR4N-TF
30773136           HMA42GR7MFR4N-TF   HMA42GR7MFR4N-TF
30772FF8           HMA42GR7MFR4N-TF   HMA42GR7MFR4N-TF
70738850           HMA42GR7MFR4N-TF   HMA42GR7MFR4N-TF
3079B2CD           HMA42GR7MFR4N-TF   HMA42GR7MFR4N-TF
307A04D2           HMA42GR7MFR4N-TF   HMA42GR7MFR4N-TF
70738951           HMA42GR7MFR4N-TF   HMA42GR7MFR4N-TF
70738939           HMA42GR7MFR4N-TF   HMA42GR7MFR4N-TF
70738A0B           HMA42GR7MFR4N-TF   HMA42GR7MFR4N-TF
70738823           HMA42GR7MFR4N-TF   HMA42GR7MFR4N-TF
```

```
Cont.:

Requires Service   Service Reason
-----------------  ------------------
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
hardware_status	Hardware Status	4
dimm_id	DIMM Id	N/A
cpu	CPU	N/A
size	Size(Mb)	N/A

Field ID	Field output	Default position
speed	Speed(MHz)	N/A
configured_speed	Configured Clock Speed(MHz)	N/A
manufacturer	Manufacturer	N/A
serial	Serial	N/A
original_serial	Original Serial	N/A
part_number	Part #	N/A
original_part_number	Original Part Number	N/A
requires_service	Requires Service	N/A
service_reason	Service Reason	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing CPUs in the modules

Use the **cpu_list** command to list the central processing units (CPU) in the modules.

```
cpu_list [ module=ModuleNumber  | cpu=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
module	Limits the listing to a specific module.	N	All CPUs in all modules.
cpu	Lists only a specific CPU.	N	A specific CPU.

Example:

```
cpu_list
```

Output:

```
Component ID   Status   Currently Functioning   Hardware Status   CPU Number   Family
------------- -------- ---------------------- ----------------- ------------ --------
1:CPU:10:1     OK       yes                    OK                1            Xeon
1:CPU:11:1     OK       yes                    OK                1            Xeon
1:CPU:12:1     OK       yes                    OK                1            Xeon
1:CPU:13:1     OK       yes                    OK                1            Xeon
1:CPU:1:1      OK       yes                    OK                1            Xeon
1:CPU:2:1      OK       yes                    OK                1            Xeon
1:CPU:3:1      OK       yes                    OK                1            Xeon
1:CPU:4:1      OK       yes                    OK                1            Xeon
1:CPU:5:1      OK       yes                    OK                1            Xeon
1:CPU:6:1      OK       yes                    OK                1            Xeon
1:CPU:7:1      OK       yes                    OK                1            Xeon
1:CPU:8:1      OK       yes                    OK                1            Xeon
1:CPU:9:1      OK       yes                    OK                1            Xeon
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
hardware_status	Hardware Status	4
number	CPU Number	5
family_string	Family	6
type_string	Type	N/A
id	ID	N/A
type	Type Code	N/A
family	Family Code	N/A
model	Model Code	N/A
stepping	Stepping	N/A
max_speed	Max Speed(MHz)	N/A
current_speed	Current Speed(MHz)	N/A
status_string	Status	N/A
manufacturer	Manufacturer	N/A
version	Version	N/A
model_string	Model	N/A
signature	Signature	N/A
core_count	Cores	N/A
core_enabled	Enabled Cores	N/A
thread_count	Threads	N/A
serial	Serial	N/A
original_serial	Original Serial	N/A
part_number	Part #	N/A
original_part_number	Original Part Number	N/A
requires_service	Requires Service	N/A
service_reason	Service Reason	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed

User Category	Permission
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
-------------- -------- ---------------------- -------------------------
1:HCA:10:1     OK       yes                    CB194A - Connect-IB QSFP
1:HCA:7:1      OK       yes                    CB194A - Connect-IB QSFP
1:HCA:9:1      OK       yes                    CB194A - Connect-IB QSFP

Cont.:

Board ID        Part Number
--------------- -------------
MT_1210110019   46W0572
MT_1210110019   46W0572
MT_1210110019   46W0572
```

Field ID	Field output	Default position
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

Field ID	Field output	Default position
original_serial	Original Serial	N/A
part_number	Part Number	6
original_part_number	Original Part Number	N/A
hardware_revision	Hardware Revision	N/A
requires_service	Requires Service	N/A
service_reason	Service Reason	N/A
adapter_id	HCA Id	N/A
guid	GUID	N/A
vendor_part_id	Vendor Part ID	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing CNA adapters in the system

Use the **cna_list** command to list CNA adapters in the storage system.

```
cna_list [ module=ModuleNumber | cna=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
module	Limits the listing to a specific module.	N	All CNA adapters in all modules.
cna	Lists only a specific CNA.	N	A specific CNA.

A converged network adapter (CNA) is a single network interface card that contains both a Fibre Channel host bus adapter and a TCP/IP Ethernet NIC. It connects servers to FC-based storage area networks (SANs) and Ethernet-based local area networks (LANs).

Example:

```
cna_list
```

Output:

```
Component ID   Status   Currently Functioning   Board Description
-------------- -------- ---------------------- ------------------------------
1:CNA:10:1     OK       yes                    CX312B - ConnectX-3 Pro SFP+
1:CNA:9:1      OK       yes                    CX312B - ConnectX-3 Pro SFP+

Board ID        Part Number
--------------- --------------
MT_1200111023   MCX312B-XCCT
MT_1200111023   MCX312B-XCCT
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
version	Version	N/A
board_description	Board Description	4
original_board_description	Original Board Description	N/A
board_id	Board ID	5
original_board_id	Original Board ID	N/A
board_type	Board Type	N/A
original_board_type	Original Board Type	N/A
serial	Serial	N/A
original_serial	Original Serial	N/A
part_number	Part Number	6
original_part_number	Original Part Number	N/A
hardware_revision	Hardware Revision	N/A
requires_service	Requires Service	N/A
service_reason	Service Reason	N/A
adapter_id	HCA Id	N/A
guid	GUID	N/A
vendor_part_id	Vendor Part ID	N/A

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing data disk devices in the system

Use the **disk_list** command to list special vault device statuses.

```
disk_list [ module=ModuleNumber  | disk=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
module	Limits the listing to a specific module.	N	All data disk devices in all modules.
disk	Data disk devices for which special statuses are to be listed.	N	All data disk devices.

This command lists the statuses of data disk devices, including:

- Component generic status
- Data disk device capacity
- Model
- Serial

Example:

```
disk_list
```

Output:

```
Component ID          Status   Currently Functioning   Capacity   Target Status   Vendor
--------------------- -------- ---------------------- ---------- --------------- -------
1:disk:14:1   OK       yes                    250GB                      LENOVO-X
1:disk:14:2   OK       yes                    250GB                      LENOVO-X
1:disk:3:1    OK       yes                    250GB                      LENOVO-X
1:disk:3:2    OK       yes                    250GB                      LENOVO-X
1:disk:5:1    OK       yes                    250GB                      LENOVO-X
1:disk:5:2    OK       yes                    250GB                      LENOVO-X
1:disk:6:1    OK       yes                    250GB                      LENOVO-X
1:disk:6:2    OK       yes                    250GB                      LENOVO-X

Cont.:

Model             Serial     Firmware   FRU       Temperature   Encryption State
----------------- ---------- ---------- --------- ------------- ------------------
HUSMR1625ASS20E   0PVGJTPA   P4C9       00NA685   22            Ready
HUSMR1625ASS20E   0PVGHN6A   P4C9       00NA685   22            Ready
HUSMR1625ASS20E   0PVJ2PEA   P4C9       00NA685   21            Ready
HUSMR1625ASS20E   0PVJ251A   P4C9       00NA685   21            Ready
HUSMR1625ASS20E   0PVJ99UA   P4C9       00NA685   20            Ready
HUSMR1625ASS20E   0PVJRB5A   P4C9       00NA685   21            Ready
HUSMR1625ASS20E   0PVJS39A   P4C9       00NA685   20            Ready
HUSMR1625ASS20E   0PVJ9RAA   P4C9       00NA685   21            Ready
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Allowed
Read-only users	Allowed
Technicians	Allowed

Listing service statuses

Use the **service_list** command to list all service specific statuses.

```
service_list [ service=ComponentId ]
```

Parameters

Name	Description	Mandatory	Default
service	The service to be listed.	N	All services

This command lists the statuses that apply to services. The list includes the following information:

- Component generic status
- Service on/failed
- Comment (optional)

Example:

```
service_list
```

Output:

```
Component ID          Status   Currently Functioning   Target Status
--------------------- -------- ---------------------- ---------------
1:Data:10             OK       yes
1:Data:11             OK       yes
1:Data:7              OK       yes
1:Data:9              OK       yes
1:Data_Reduction:10   OK       yes
1:Data_Reduction:11   OK       yes
1:Data_Reduction:7    OK       yes
1:Data_Reduction:9    OK       yes
1:Interface:10        OK       yes
1:Interface:11        OK       yes
1:Interface:9         OK       yes
1:Remote:10           OK       yes
1:Remote:11           OK       yes
1:Remote:9            OK       yes
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3
target_status	Target Status	4

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed

User Category	Permission
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
```

Field ID	Field output	Default position
component_id	Component ID	1
status	Status	2
currently_functioning	Currently Functioning	3

Field ID	Field output	Default position
requires_service	Requires Service	4
service_reason	Service Reason	5

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Listing trace snapshot on a module

Use the **traces_snapshot_list** command to list trace snapshots on a module.

```
traces_snapshot_list module=ModuleNumber
```

Parameters

Name	Description	Mandatory
module	Component ID of the module to query.	Y

Field ID	Field output	Default position
snapshot	Snapshot Directories	1

Example:

```
traces_snapshot_list  module=1:Module:9
```

Output:

```
Snapshots Directories
---------------------------------
1__20120802_1653_20120802_1713
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Allowed

Creating a trace snapshot

Use the **traces_snapshot** command to create a trace snapshot.

```
traces_snapshot [ snapshot_back_time=MINUTES ] [ snapshot_delay_time=MINUTES ]
```

Parameters

Name	Type	Description	Mandatory	Default
snapshot_delay_time	Integer	Max delay between the request and snapshot creation.	N	no. Uses configuration misc.internal.auto_snapshot_trace.last_snapshot_minutes_delay field.
snapshot_back_time	Integer	Time back from the request time to include in the snapshot.	N	no. Uses configuration misc.internal.auto_snapshot_trace.snapshot_back_time field.

Example:

```
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

Chapter 22. 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.

```
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.

© Copyright IBM Corp. 2017

Name	Type	Description	Mandatory	Default
start	N/A	Starting point for the statistics report.	N	N/A
end	N/A	Ending point for the statistics report.	N	N/A
count	Positive integer	Number of time points reported.	Y	N/A
interval	Positive integer	The length of time in each statistic's time point. The resolution of this number is set in *resolution_unit*.	Y	N/A
resolution_unit	Enumeration	Sets the unit of measurement for the length of each bin.	Y	N/A
perf_class	Object name	Displays performance class aggregated statistics for bandwidth and IOPS.	N	All Performance classes.

This command lists I/O statistics. The **count** parameter sets the number of lines in the statistics report. The combination of the **interval** and **resolution_unit** parameters sets the length of time for each statistics line. Either start or end timestamp must be provided. These timestamps set the time for the statistics report. Other parameters restrict statistics to a specific host, host port, volume, domain, interface port and so on.

For each line of statistics, 48 numbers are reported, which represent all the combinations of reads/writes, hits/misses and I/O size reporting for each of the 16 options for bandwidth, IOPS and latency. Statistics collection is limited to 32 pools and 200 volumes.

The syntax for the **start** and **end** fields is as follows: `Y-M-D[.[h[:m[:s]]]]`, where the ranges are as follows:

- Y - year (four digit)
- M - month (1-12)
- D - day (1-31)
- h - hour (0-23, with 0 as default)
- m - minute (0-59, with 0 as default)
- s - second (0-59, with 0 as default)

The year, month and day are separated by dashes, and the optional hours, minutes and seconds are separated by colons.

Output units:

- Very Large blocks are >512KB
- Large blocks - 64-512KB
- Medium blocks - 8-64KB
- Small blocks - 0-8KB

- The latency is in Microseconds
- The bandwidth is in KB

Field ID	Field output	Default position
time	Time	1
failures	Failures	N/A
aborts	Aborts	N/A
read_hit_very_large_iops	Read Hit Very large - IOps	2
read_hit_very_large_latency	Read Hit Very large - Latency	3
read_hit_very_large_throughput	Read Hit Very large - Throughput	4
read_hit_very_large_remotely_served	Read Hit Very large - Remotely Served IOs	63
read_hit_large_iops	Read Hit Large - IOps	5
read_hit_large_latency	Read Hit Large - Latency	6
read_hit_large_throughput	Read Hit Large - Throughput	7
read_hit_large_remotely_served	Read Hit Large - Remotely Served IOs	64
read_hit_medium_iops	Read Hit Medium - IOps	8
read_hit_medium_latency	Read Hit Medium - Latency	9
read_hit_medium_throughput	Read Hit Medium - Throughput	10
read_hit_medium_remotely_served	Read Hit Medium - Remotely Served IOs	65
read_hit_small_iops	Read Hit Small - IOps	11
read_hit_small_latency	Read Hit Small - Latency	12
read_hit_small_throughput	Read Hit Small - Throughput	13
read_hit_small_remotely_served	Read Hit Small - Remotely Served IOs	66
read_miss_very_large_iops	Read Miss Very large - IOps	14
read_miss_very_large_latency	Read Miss Very large - Latency	15
read_miss_very_large_throughput	Read Miss Very large - Throughput	16
read_miss_very_large_remotely_served	Read Miss Very large - Remotely Served IOs	67
read_miss_large_iops	Read Miss Large - IOps	17
read_miss_large_latency	Read Miss Large - Latency	18
read_miss_large_throughput	Read Miss Large - Throughput	19
read_miss_large_remotely_served	Read Miss Large - Remotely Served IOs	68
read_miss_medium_iops	Read Miss Medium - IOps	20
read_miss_medium_latency	Read Miss Medium - Latency	21
read_miss_medium_throughput	Read Miss Medium - Throughput	22
read_miss_medium_remotely_served	Read Miss Medium - Remotely Served IOs	69
read_miss_small_iops	Read Miss Small - IOps	23
read_miss_small_latency	Read Miss Small - Latency	24
read_miss_small_throughput	Read Miss Small - Throughput	25
read_miss_small_remotely_served	Read Miss Small - Remotely Served IOs	70
write_hit_very_large_iops	Write Hit Very large - IOps	26

Field ID	Field output	Default position
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_very_large_ remotely_served	Read Memory-Hit Very large - Remotely Served IOs	71
read_memory_hit_large_iops	Read Memory-Hit Large - IOps	53
read_memory_hit_large_latency	Read Memory-Hit Large - Latency	54
read_memory_hit_large_ throughput	Read Memory-Hit Large - Throughput	55
read_memory_hit_large_ remotely_served	Read Memory-Hit Large - Remotely Served IOs	72
read_memory_hit_medium_iops	Read Memory-Hit Medium - IOps	56
read_memory_hit_medium_latency	Read Memory-Hit Medium - Latency	57
read_memory_hit_medium_ throughput	Read Memory-Hit Medium - Throughput	58
read_memory_hit_medium_ remotely_served	Read Memory-Hit Medium - Remotely Served IOs	73
read_memory_hit_small_iops	Read Memory-Hit Small - IOps	59

Field ID	Field output	Default position
read_memory_hit_small_latency	Read Memory-Hit Small - Latency	60
read_memory_hit_small_ throughput	Read Memory-Hit Small - Throughput	61
read_memory_hit_small_ remotely_served	Read Memory-Hit Small - Remotely Served IOs	74
time_in_seconds	Time (s)	62

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Return codes

- **BAD_TIME_FORMAT**
  Bad time format. Should be YYYY-MM-DD[.HH[:MM[:SS]]]
- **TARGET_PORT_BAD_ADDRESS**
  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
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

Chapter 23. 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.

```
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

© Copyright IBM Corp. 2017

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, smsgw, smtpgw, target, user, user_group, vol.	Y
name	Object name	The name of the target object.	Y
key	String	Metadata key.	Y

This command deletes a metadata key value for the specified object.

The command will fail if the key is not defined.

Example:

```
metadata_delete object_type=host name=Host1 key=01
```

Output:

```
Command completed successfully.
```

Access control

User Category	Permission	Condition
Storage administrator	Allowed	N/A
Storage integration administrator	Allowed	N/A
Application administrator	Conditionally Allowed	Metadata can be set for only volumes, snapshots, snapshot groups, clusters or hosts, and only for objects associated with the application administrator executing the command. Hosts or clusters should be associated with the user. Volumes should be mapped to a host or a cluster associated with the user. Snapshots or snapshot groups should be ones created by application administrator.
Security administrator	Disallowed	N/A
Read-only users	Disallowed	N/A
Technicians	Disallowed	N/A

Return codes

- **OBJECT_BAD_NAME**

  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, smsgw, smtpgw, target, user, user_group, vol.
name	Object name	The name of the target object.	N	All objects
key	String	Metadata key.	N	List all keys and values.
domain	Object name	The domain name.	N	All Domains

This command lists all the value key pairs for this object, or a specific one. The command fails if no key is defined.

Example:

```
metadata_list object_type=host
```

Output:

```
Object Type   Name    Key   Value
------------- ------- ----- -------
Host          Host1   01    Host
```

Field ID	Field output	Default position
object_type	Object Type	1
name	Name	2
key	Key	3
value	Value	4

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Allowed
Application administrator	Allowed
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Allowed

Setting user metadata

Use the **user_metadata_set** command to set user metadata by section and key.

```
user_metadata_set  section=Section key=Key value=Value
```

Parameters

Name	Type	Description	Mandatory
section	Enumeration	Metadata section.	Y
key	String	Metadata key.	Y
value	String	Metadata value.	Y

Example:

```
user_metadata_set section=GUI key=GUI_10.2 value=Ubunto_2.6
```

Output:

```
(Command returns no output)
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- **USER_METADATA_FULL**

  Maximal number of user metadata objects has been reached.

Listing user metadata

Use the **user_metadata_list** command to list user metadata.

```
user_metadata_list [ section=Section ]
```

Parameters

Name	Type	Description	Mandatory	Default
section	Enumeration	User metadata section.	N	''

Example:

```
user_metadata_list
```

Output:

```
Time                 Section  Key       Value
-------------------- -------- --------- -------
2004-11-22 18:08:23   TA       GUI_0.2   Ubunto_2.2
2007-11-22 18:08:22   GUI      GUI_10.1  Ubunto_2.6
2007-11-22 18:08:23   GUI      GUI_10.2  Ubunto_2.6
```

Field ID	Field output	Default position
time	Time	1
section	Section	2
key	Key	3
value	Value	4

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Allowed

User Category	Permission
Security administrator	Disallowed
Read-only users	Allowed
Technicians	Disallowed

Deleting user metadata

Use the **user_metadata_delete** command to delete user metadata by user section and key.

```
user_metadata_delete section=Section key=Key
```

Parameters

Name	Type	Description	Mandatory
section	Enumeration	Meta data section.	Y
key	String	Metadata key.	Y

Example:

```
user_metadata_delete section=GUI key=GUI_10.2
```

Output:

```
(Command returns no output)
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Disallowed
Read-only users	Disallowed
Technicians	Disallowed

Return codes

- **USER_METADATA_OBJECT_KEY_NOT_FOUND**

  The specified user metadata object does not exist.

Chapter 24. 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.

```
encrypt_disable
```

This command disables the data protection feature.

In order for this command to complete successfully, all of the following prerequisites must be fulfilled:

- The system is fully redundant
- None of the present and active SSDs, Flash cards, or Flash canisters is failed
- No Flash enclosure is undergoing CCL

When data protection is disabled, a cryptographic erase is performed on all protected bands (ensuring that all existing user data is no longer accessible). After the command successfully completes, all bands are left in the unlocked state.

Disabling encryption when the encryption state is other than Active (displayed as `Enabled` in **state_list**) will result in an error.

Example:

```
encrypt_disable -y
```

Output:

```
Command 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

© Copyright IBM Corp. 2017

Return codes

- **CANNOT_DISABLE_ENCRYPTION_WHILE_NOT_IN_FULLY_PROTECTED_STATE**

  Cannot disable encryption while not in fully protected state.

  Troubleshooting: Resolve any issues preventing system from reaching full data protection state. Contact IBM support for assistance.
- **UNSUPPORTED_HARDWARE**

  Cannot utilize encryption on unsupported hardware.

  Troubleshooting: Contact IBM 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: Contact IBM Support
- **DATA_REDUCTION_RECOVERY_IS_RUNNING**

  Data reduction recovery is running, operation is not allowed

  Troubleshooting: Wait for data reduction recovery to complete.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  Data reduced tier is offline, operation is not allowed

  Troubleshooting: Contact IBM Support
- **SECURE_ERASE_IS_RUNNING**

  Operation not allowed while secure erase is running.

  Troubleshooting: Wait for secure erase process to complete.
- **FLASH_CCL_IN_PROGRESS**

  The requested command cannot be invoked while Flash Enclosure CCL is in progress.

  Troubleshooting: Wait for Flash CCL to complete
- **ENCRYPTION_CANNOT_UNENROLL_SOME_VAULT_DEVICES**

  Failed components prevent unenrolling some vault devices.

  Troubleshooting: Contact IBM Support
- **ENCRYPTION_CANNOT_UNENROLL_SOME_FLASH_ENCLOSURES**

  Not all flash enclosures are phased in before unenrollment.

  Troubleshooting: Contact IBM Support
- **ENCRYPTION_IS_NOT_IN_A_STABLE_STATE**

  Encryption is not in a stable state.

  Troubleshooting: Wait for encryption process to complete.

Enabling encryption

Use the **encrypt_enable** command to enable the data protection feature.

```
encrypt_enable [ recovery_keys=<yes|no> ] [ key_scheme=KeyScheme ]
```

Parameters

Name	Type	Description	Mandatory	Default
recovery_keys	Boolean	Defines whether recovery keys are required for encryption activation.	N	yes
key_scheme	Enumeration	Defines which key scheme to use for encryption activation: external or local.	N	external

This command enables the data protection feature. In order for this command to complete successfully, all of the following prerequisites must be fulfilled:

- The system is fully redundant
- None of the present and active SSDs, Flash cards, or Flash canisters is failed
- No Flash enclosure is undergoing CCL

Example:

```
encrypt_enable recovery_keys=yes key_scheme=local -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
- **ARE_YOU_SURE_YOU_WANT_TO_ENABLE_LOCAL_KEY_MANAGEMENT_ENCRYPTION**

  Are you sure you want to enable encryption with local key management on this system?

  Troubleshooting: A yes option is required for this command

Return codes

- **CANNOT_ENABLE_ENCRYPTION_WHILE_NOT_IN_FULLY_PROTECTED_STATE**

  Cannot enable encryption while not in fully protected state.

Troubleshooting: Resolve any issues preventing system from reaching full data protection state. Contact IBM support for assistance.

- **UNSUPPORTED_HARDWARE**

  Cannot utilize encryption on unsupported hardware.

  Troubleshooting: Contact IBM 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
- **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.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  Data reduced tier is offline, operation is not allowed

  Troubleshooting: Contact IBM Support
- **ENCRYPTION_CANNOT_ENROLL_SOME_FLASH_ENCLOSURES**

  Not all flash enclosures are phased in before enrollment.

  Troubleshooting: Contact IBM Support
- **ENCRYPTION_CANNOT_ENROLL_SOME_VAULT_DEVICES**

  Failed components prevent enrolling some vault devices.

  Troubleshooting: Contact IBM Support
- **DATA_REDUCTION_RECOVERY_IS_RUNNING**

  Data reduction recovery is running, operation is not allowed

  Troubleshooting: Wait for data reduction recovery to complete.
- **FLASH_CCL_IN_PROGRESS**

  The requested command cannot be invoked while Flash Enclosure CCL is in progress.

  Troubleshooting: Wait for Flash CCL to complete
- **RECOVERY_MISMATCH_KEY_SCHEME**

  The recovery key scheme mismatch to the current scheme.

  Troubleshooting: Check the recovery key scheme using encrypt_key_scheme_get.

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---*MIICyTCCAbGgAwIBAgIXLSiyd2FPMA0GCSqGSIb3IiEBCwUAMBQx
EjAQAgNVBVuTCXNrbG5pdHNv*......
*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
- **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 keyservers currently defined in the system along with their connectivity status.

```
encrypt_keyserver_list
```

Example:

```
encrypt_keyserver_list
```

Output:

```
Module   Name      App/Key Status   Last time checked   Master   Port
3        nachos    NOAPP            2013/03/27 20:18:43 yes      5696
3        nachos    UNKNOWN          2013/03/27 20:18:43 yes      5696
3        snocone   UNKNOWN          2013/03/27 20:18:43 no       5696
3        snocone   ACTIVE           2013/03/27 20:18:43 no       5696
3        TKLM-SA   BAD_CERT         2013/03/27 20:18:43 no       5696

Address
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.

Checking keyserver status

Use the **encrypt_keyserver_check_status** command to check connectivity status of the keyservers currently defined in the system.

```
encrypt_keyserver_check_status
```

This command initiates the async check of the connectivity status of the keyservers currently defined in the system. To get the current status, use the **encrypt_keyserver_list** CLI command.

Example:

```
encrypt_keyserver_check_status
```

Output:

```
Command completed successfully
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Allowed
Technicians	Allowed

Return codes

- **UNSUPPORTED_HARDWARE**

  Cannot 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
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
- **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 keyservers. To unlock the disks, the **min_req number** (defined by the **encrypt_recovery_key_generate** command) of security administrators must all successfully enter their recovery key (as presented to them via **recovery_key_get**). After the minimum required keys have been entered, the storage administrator must change the state from Maintenance to On by issuing

state_change target_state=on. When this command is issued with the machine in the On state, it has no effect, and can be used to check the validity of the recovery key.

Example:

```
encrypt_recovery_key_enter
key=CBC9B398373FDE79CD38B23192DABACADB5DA63A915CBF5CA8C4E0C212819DE6
```

Output:

```
Command 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 IBM 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 IBM 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 IBM 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 IBM 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 ] [ key_scheme=KeyScheme ]
```

Parameters

Name	Type	Description	Mandatory	Default
min_req	Integer	Minimum number of required security administrator recovery key shares.	N	2
users	Object name	User names of the security administrators.	Y	N/A
key_scheme	Enumeration	Defines which key scheme to use for encryption activation external or local.	N	external

This command is used to specify which security administrator will receive recovery keys (or, more accurately, "shares"), and to define the minimum number of recovery keys that need to be entered (using the **encrypt_recovery_key_enter** command) in order to unlock the encrypted keys. Once this command has been entered, all the specified security administrators are expected to retrieve and verify their recovery keys, using **encrypt_recovery_key_get** and **encrypt_recovery_key_verify**, respectively. This command can only be run when **encryption_state** is DISABLED.

Example:

```
xcli -u secadmin1 -p password -m ${HOST} 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 IBM 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 IBM 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 IBM 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 IBM Support
- **ENCRYPTION_KR_READ_FAILED**

  Error reading the key repository.

  Troubleshooting: Contact IBM Support
- **CANNOT_UPDATE_KEY_METADATA**

  Cannot update metadata in key repository for new key.

  Troubleshooting: Contact IBM Support

Retrieving the security administrator's recovery key

Use the **encrypt_recovery_key_get** command to retrieve the recovery key share generated for the current user.

```
encrypt_recovery_key_get
```

This command retrieves the recovery key generated for the current user (by issuing **encrypt_recovery_key_generate** or **encrypt_recovery_key_rekey**) to be stored in a secure manner. After running this command, the user needs to 'prove' that they have the key by entering it via the **encrypt_recovery_key_verify** command. Once this is completed successfully, **encrypt_recovery_key_get** will no longer present the user's key. Using **encrypt_recovery_key_get** more than once will return the same value again.

Example:

```
encrypt_recovery_key_get
```

Output:

```
Command 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

Name	Type	Description	Mandatory	Default
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
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
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, 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.
- **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

Obtaining a new master key

Use the **encrypt_local_rekey** command to initiate rekeying (getting new cryptographic material) from the local key management.

```
encrypt_local_rekey
```

Example:

```
encrypt_local_rekey
```

Output:

```
Command 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_REKEY**

  Are you sure you want to change the local key?

  Troubleshooting: A yes option is required for this command

Return codes

- **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 service is active and serving keys (obtained from key server or stored local key). Otherwise, contact IBM 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 IBM Support
- **CANNOT_CANNOT_GENERATE_EXMK_ESKH**

  Problem generating EXMK and ESKH.

  Troubleshooting: Contact IBM Support
- **CANNOT_WRITE_TO_KEY_REPOSITORY**

  Failed writing keys to the key repository.

  Troubleshooting: Contact IBM Support
- **CANNOT_COPY_KEYS_IN_KEY_REPOSITORY**

  Problem copying current keys to old keys location in the key repository.

  Troubleshooting: Contact IBM Support
- **ENCRYPTION_KR_WRITE_FAILED**

  Error writing to the key repository.

Troubleshooting: Contact IBM Support

- **ENCRYPTION_KR_READ_FAILED**
  Error reading the key repository.
  Troubleshooting: Contact IBM Support
- **NO_LIVE_KEYSERVER_GATEWAY_NODE**
  There is no live key server gateway node on the system.
  Troubleshooting: Please restart the key server gateway node and try again.
- **ENCRYPTION_NOT_IN_LOCAL_SCHEME**
  Encryption is not in local key scheme.
  Troubleshooting: Check encrypt_key_scheme_get command

Changing the key management scheme

Use the **encrypt_change_key_scheme** command to change the key management scheme.

```
encrypt_change_key_scheme key_scheme=KeyScheme
```

Parameters

Name	Type	Description	Mandatory
key_scheme	Enumeration	Defines which key management scheme (external or local) to use for encryption activation.	Y

This command only supports the change from an external to a local scheme.

Example:

```
encrypt_change_key_scheme key_scheme=local
```

Output:

```
Command 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_CHANGE_KEY_SCHEME**

Are you sure you want to change the key management scheme? This operation is irreversible.

Troubleshooting: A yes option is required for this command

Return codes

- **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 service is active and serving keys (obtained from key server or stored local key). Otherwise, contact IBM 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 IBM Support
- **CANNOT_CANNOT_GENERATE_EXMK_ESKH**

  Problem generating EXMK and ESKH.

  Troubleshooting: Contact IBM Support
- **CANNOT_WRITE_TO_KEY_REPOSITORY**

  Failed writing keys to the key repository.

  Troubleshooting: Contact IBM Support
- **CANNOT_COPY_KEYS_IN_KEY_REPOSITORY**

  Problem copying current keys to old keys location in the key repository.

  Troubleshooting: Contact IBM Support
- **ENCRYPTION_KR_WRITE_FAILED**

  Error writing to the key repository.

  Troubleshooting: Contact IBM Support
- **ENCRYPTION_KR_READ_FAILED**

  Error reading the key repository.

  Troubleshooting: Contact IBM Support
- **NO_LIVE_KEYSERVER_GATEWAY_NODE**

  There is no live key server gateway node on the system.

  Troubleshooting: Please restart the key server gateway node and try again.
- **ENCRYPTION_NOT_IN_EXTERNAL_SCHEME**

  Encryption is not in external key scheme.

  Troubleshooting: Check encrypt_key_scheme_get command

Viewing the key scheme

Use the **encrypt_key_scheme_get** command to view the key scheme defined in the system.

```
encrypt_key_scheme_get
```

Example:

```
encrypt_key_scheme_get
```

Output:

```
Command executed successfully.
    encrypt_key_scheme = "LOCAL"
```

Access control

User Category	Permission
Storage administrator	Allowed
Storage integration administrator	Disallowed
Application administrator	Disallowed
Security administrator	Allowed
Read-only users	Allowed
Technicians	Allowed

Chapter 25. Events

VOLUME_CREATE

Severity	Description
informational	Volume was created with name '*volume.name*' and size *volume.size*GB in Storage Pool with name '*volume.pool_name*'.

VOLUME_CREATE_FAILED_TOO_MANY

Severity	Description	Troubleshooting
warning	Volume with name '*name*' could not be created. You are attempting to add more volumes than the system permits.	Delete volumes to allow new ones to be created.

VOLUME_RENAME

Severity	Description
informational	Volume with name '*old_name*' was renamed '*volume.name*'.

SUBORDINATE_VOLUME_RENAME

Severity	Description
informational	Remote volume with name '*old_name*' was renamed '*volume.name* by local system'.

VOLUME_RESIZE

Severity	Description
informational	Volume with name '*volume.name*' was resized from *old_size*GB to *volume.size*GB.

SECONDARY_VOLUME_RESIZE

Severity	Description
informational	Secondary volume with name '*volume.name*' was resized by primary machine from *old_size*GB to *volume.size*GB.

VOLUME_DELETE

Severity	Description
informational	Volume with name '*volume.name*' was deleted and its data is no longer accessible.

© Copyright IBM Corp. 2017

VOLUME_FORMAT

Severity	Description
informational	Volume with name '*volume.name*' was formatted.

VOLUME_COPY

Severity	Description
informational	Volume with name '*source.name*' was copied to volume with name '*target.name*'.

VOLUME_COPY_DIFF

Severity	Description
informational	Volume with name '*source.name*' was diff-copied from base '*base.name*' to volume with name '*target.name*'.

VOLUME_LOCK

Severity	Description
informational	Volume with name '*volume.name*' was locked and set to 'read-only'.

VOLUME_UNLOCK

Severity	Description
informational	Volume with name '*volume.name*' was unlocked and set to 'writable'.

SUBORDINATE_VOL_LOCK

Severity	Description
informational	Remote volume with name '*volume.name*' was locked and set to 'read-only' by local machine'.

SUBORDINATE_VOL_UNLOCK

Severity	Description
informational	Remote volume with name '*volume.name*' was unlocked and set to 'writable' by local machine.

VOLUME_MOVE

Severity	Description
informational	Volume with name '*volume.name*' has been moved from Storage Pool '*orig_pool.name*' to Pool '*pool.name*'.

OLVM_CREATE

Severity	Description
informational	Volume with name '*volume.name*' was defined as an IBM Hyper-Scale Mobility.

OLVM_OWNER_CREATE

Severity	Description
informational	Volume with name '*volume.name*' was defined as an OWNER IBM Hyper-Scale Mobility.

OLVM_OWNER_ACTIVATED

Severity	Description
informational	IBM Hyper-Scale Mobility Owner Volume with name '*volume.name*' was activated.

OLVM_ACTIVATE

Severity	Description
informational	IBM Hyper-Scale Mobility Volume with name '*volume.name*' was activated.

OLVM_DEACTIVATE

Severity	Description
informational	IBM Hyper-Scale Mobility Volume with name '*volume.name*' was deactivated.

OLVM_REMOTE_ACTIVATE

Severity	Description
informational	IBM Hyper-Scale Mobility Volume with name '*volume.name*' was activated.

OLVM_REMOTE_DEACTIVATE

Severity	Description
informational	IBM Hyper-Scale Mobility Volume with name '*volume.name*' was deactivated.

OLVM_PROXY_INITIATED

Severity	Description
informational	IBM Hyper-Scale Mobility Volume with name '*volume.name*' move to Proxy state started.

OLVM_PROXY

Severity	Description
informational	IBM Hyper-Scale Mobility Volume with name '*volume.name*' entered Proxy state. Volume data on the system is freed.

OLVM_DELETE

Severity	Description	Troubleshooting
informational	IBM Hyper-Scale Mobility Volume process with name '*name*' was deleted.	Delete volumes to allow new ones to be created.

OLVM_ABORT

Severity	Description
informational	IBM Hyper-Scale Mobility Volume process with name '*volume.name*' was aborted.

OLVM_OWNER_DELETE

Severity	Description
informational	IBM Hyper-Scale Mobility Owner Volume process with name '*volume.name*' was deleted.

OLVM_OWNER_ABORT

Severity	Description
informational	IBM Hyper-Scale Mobility Owner Volume process with name '*volume.name*' was aborted.

OLVM_SYNC_STARTED

Severity	Description
informational	IBM Hyper-Scale Mobility Synchronization of volume '*name*' has started.

OLVM_SYNC_ENDED

Severity	Description
informational	IBM Hyper-Scale Mobility Synchronization of volume '*name*' has ended.

QUORUM_WITNESS_DEFINED

Severity	Description
informational	Quorum witness '*Quorum Witness Name*' defined.

QUORUM_WITNESS_DELETED

Severity	Description
informational	Quorum witness '*Quorum Witness Name*' deleted.

QUORUM_WITNESS_UPDATED

Severity	Description
informational	Quorum witness '*Quorum Witness Name*' updated.

QUORUM_WITNESS_RENAMED

Severity	Description
informational	Quorum witness with name '*Old Name*' was renamed '*New Name*'.

QUORUM_WITNESS_ACTIVATION_START

Severity	Description
informational	Quorum witness '*Quorum Witness Name*' activation started.

QUORUM_WITNESS_ACTIVATION_SUCCESSFUL

Severity	Description
informational	Quorum witness '*Quorum Witness Name*' activated successfully.

QUORUM_WITNESS_ACTIVATION_FAILED

Severity	Description
minor	Quorum witness '*Quorum Witness Name*' activation failed.

QUORUM_WITNESS_DEACTIVATION_START

Severity	Description
informational	Quorum witness '*Quorum Witness Name*' deactivation started.

QUORUM_WITNESS_DEACTIVATION_SUCCESSFUL

Severity	Description
informational	Quorum witness '*Quorum Witness Name*' deactivated successfully.

QUORUM_WITNESS_DEACTIVATION_FAILED

Severity	Description
minor	Quorum witness '*Quorum Witness Name*' deactivation failed.

QUORUM_WITNESS_COMMUNICATION_DOWN

Severity	Description
major	Communication with quorum witness '*Quorum Witness Name*' is down.

QUORUM_WITNESS_COMMUNICATION_UP

Severity	Description
informational	Communication with quorum witness '*Quorum Witness Name*' is up.

QUORUM_WITNESS_HEARTBEATS_OPERATIONAL

Severity	Description
informational	Successfully sending heartbeats to quorum witness '*Quorum Witness Name*'.

QUORUM_WITNESS_HEARTBEATS_FAILING

Severity	Description
major	Failing to send heartbeats to quorum witness '*Quorum Witness Name*'.

QUORUM_WITNESS_CERTIFICATE_ABOUT_TO_EXPIRE

Severity	Description
variable	Quorum witness '*Quorum Witness Name*' certificate is about to expire on *Expiration Date* (*Counter* notification).

QUORUM_WITNESS_CERTIFICATE_HAS_EXPIRED

Severity	Description
critical	Quorum witness '*Quorum Witness Name*' certificate has expired.

QUORUM_WITNESS_EXTERNAL_NAME_CHANGED

Severity	Description
informational	Quorum witness '*Quorum Witness Name*' has acquired a new external name '*External Name*'.

QUORUM_WITNESS_EVENT_REPORT

Severity	Description
variable	Event reported from quorum witness '*Quorum Witness Name*': *Event Description* (*System ID*).

QUORUM_WITNESS_MISSING_EVENTS

Severity	Description
informational	The events from '*First*' through '*Last*' are missing from quorum witness '*Quorum Witness Name*'.

QUORUM_WITNESS_DB_RECOVERY_NEEDED

Severity	Description
critical	Quorum witness '*Quorum Witness Name*' db has failed. Recovery is needed.

QUORUM_WITNESS_DB_RESTORE_PENDING

Severity	Description
critical	Quorum witness '*Quorum Witness Name*' db restore is pending. The db will be restored when the quorum witness is activated.

QUORUM_WITNESS_DB_RECOVERED

Severity	Description
major	Quorum witness '*Quorum Witness Name*' db was successfully recovered.

QUORUM_WITNESS_LOG_AUTHENTICATION_SET

Severity	Description
informational	Quorum witness '*Quorum Witness Name*' authentication for log retrieval was set.

QUORUM_WITNESS_LOG_AUTHENTICATION_SET_FAILED

Severity	Description
minor	Quorum witness '*Quorum Witness Name*' authentication for log retrieval setup failed, '*Failure Reason*'.

QUORUM_WITNESS_LOG_AUTHENTICATION_RESET

Severity	Description
informational	Quorum witness '*Quorum Witness Name*' authentication for log retrieval was reset.

QUORUM_WITNESS_LOG_AUTHENTICATION_RESET_FAILED

Severity	Description
minor	Quorum witness '*Quorum Witness Name*' authentication reset for log retrieval setup failed, '*Failure Reason*'.

HA_CREATE

Severity	Description
informational	A HA relation was defined for Volume '*local volume name*' on Target '*target name*'.

HA_CREATE_SLAVE

Severity	Description
informational	A HA relation was defined by Target '*target name*' for Volume '*local volume name*'.

HA_ACTIVATE

Severity	Description
informational	The HA relation of peer '*local peer name*' on Target '*target name*' was activated.

HA_SLAVE_ACTIVATE

Severity	Description
informational	The HA relation on peer '*local peer name*' was activated by target '*target name*'.

HA_SLAVE_DEACTIVATE

Severity	Description
informational	The HA relation on peer '*local peer name*' was deactivated by target '*target name*'.

HA_INCOMPATIBLE_VERSION_FOR_UNMAP_SUPPORT

Severity	Description
warning	A HA of peer '*local peer name*' on target '*target name*' cannot support unmap, remote machine has incompatible version.

CG_HA_CREATE

Severity	Description
informational	High availability mirror was defined for Consistency Group '*local CG name*' on Target '*target name*'. Remote Consistency Group is '*remote CG name*'.

CG_HA_CREATE_SLAVE

Severity	Description
informational	High availability mirror was defined by Target '*target name*' for CG '*local CG name*'. Remote CG is '*remote CG name*'.

HA_SYNC_STARTED

Severity	Description
informational	Synchronization of remote ha relation of volume '*local volume name*' on Target '*target name*' has started.

HA_SYNC_ENDED

Severity	Description
informational	Synchronization of remote ha relation of peer '*local peer name*' on target '*target name*' has ended.

HA_DEACTIVATE

Severity	Description
informational	The HA relation of peer '*local peer name*' on Target '*target name*' was deactivated.

HA_DELETE

Severity	Description
informational	The HA relation of peer '*local peer name*' to a peer on Target '*target name*' was deleted.

HA_SWITCH_ROLES_TO_SLAVE

Severity	Description
informational	Local peer '*local peer name*' is now Slave of a peer on Target '*target name*'.

HA_SWITCH_ROLES_TO_MASTER

Severity	Description
informational	Local peer '*local peer name*' is now Master of a peer on Target '*target name*'.

HA_REVERSE_ROLE_TO_SLAVE

Severity	Description
informational	Local peer '*local peer name*' is now Slave of a peer on Target '*target name*'.

HA_REVERSE_ROLE_TO_MASTER

Severity	Description
informational	Local peer '*local peer name*' is now Master of a peer on Target '*target name*'.

HA_AVAILABILITY_RESTORED

Severity	Description
informational	Local peer '*local peer name*' was made available.

HA_SNAPSHOT_CREATE

Severity	Description
informational	HA Snapshot named '*snapshot.name*' was created for volume named '*volume.name*'.

HA_SNAPSHOT_CREATE_FAILED

Severity	Description
minor	HA Remote snapshot named '*snapshot name*' was not created successfully. Error code is '*error*'

HA_SNAPSHOT_OVERWRITE

Severity	Description
informational	HA Snapshot named '*snapshot.name*' was overwritten for volume named '*volume.name*'.

HA_SLAVE_SNAPSHOT_CREATE

Severity	Description
informational	HA Snapshot named '*snapshot.name*' was created for volume named '*volume.name*'.

HA_SLAVE_SNAPSHOT_OVERWRITE

Severity	Description
informational	HA Snapshot named '*snapshot.name*' was overwritten for volume named '*volume.name*'.

HA_HIGH_AVAILABILITY_ENABLED

Severity	Description
informational	HA relation on peer '*local peer name*' high availability is enabled.

HA_HIGH_AVAILABILITY_DISABLED

Severity	Description
informational	HA relation on peer '*local peer name*' high availability is disabled by user.

HA_CONVERTED_INTO_MIRROR

Severity	Description
informational	HA relation of peer '*local peer name*' on Target '*target name*' was converted into mirror.

HA_CONVERTED_INTO_MIRROR_SLAVE

Severity	Description
informational	HA relation of peer '*local peer name*' on Target '*target name*' was converted into mirror.

TARGET_CONNECTION_HA_SUFFICIENT

Severity	Description
informational	Target named '*target.name*' is ha sufficient connected.

TARGET_CONNECTION_HA_INSUFFICIENT

Severity	Description
critical	Target named '*target.name*' is ha insufficient connected.

TARGET_IS_HA_HEALTHY

Severity	Description
informational	Target named '*target.name*' is HA healthy according to quorum witness.

TARGET_IS_HA_UNHEALTHY

Severity	Description
major	Target named '*target.name*' is HA unhealthy according to quorum witness.

TARGET_DATA_SERVICE_FAILURE

Severity	Description
major	Target named '*target.name*' has data service failure. Reason: *Reason*.

TARGET_RESUMED_NORMAL_OPERATION

Severity	Description
informational	Target named '*target.name*' resumed normal operation.

HA_AUTOMATIC_FAILOVER_SUCCESSFUL

Severity	Description
warning	HA Slave relation on peer '*local peer name*' has completed failover.

HA_MASTER_REMAINS_AVAILABLE

Severity	Description
informational	HA Master relation on peer '*local peer name*' remains available. Reason: *Reason*.

HA_MASTER_AVAILABLE

Severity	Description
informational	HA Master relation on peer '*local peer name*' is available.

HA_MASTER_UNAVAILABLE

Severity	Description
major	HA Master relation on peer '*local peer name*' is unavailable. Reason: *Unavailable Reason*.

HA_SLAVE_AVAILABLE

Severity	Description
informational	HA Slave relation on peer '*local peer name*' is available.

HA_SLAVE_UNAVAILABLE

Severity	Description
major	HA Slave relation on peer '*local peer name*' is unavailable. Reason: *Unavailable Reason*.

HA_MASTER_RELEASED_CONTROL

Severity	Description
informational	HA Master released control on relation '*local peer name*'.

HA_SLAVE_RELEASED_CONTROL

Severity	Description
informational	HA Slave released control on relation '*local peer name*'.

MIRROR_CONVERT_INTO_HA_STARTED

Severity	Description
informational	Remote mirror of peer '*local peer name*' on Target '*target name*' conversion into HA started.

MIRROR_CONVERT_INTO_HA_STARTED_SLAVE

Severity	Description
informational	Remote mirror of peer '*local peer name*' on Target '*target name*' conversion into HA started.

MIRROR_CONVERT_INTO_HA_ENDED

Severity	Description
informational	Remote mirror of peer '*local peer name*' on Target '*target name*' conversion into HA ended.

MIRROR_CONVERT_INTO_HA_ENDED_SLAVE

Severity	Description
informational	Remote mirror of peer '*local peer name*' on Target '*target name*' conversion into HA ended.

HA_REESTABLISH_FAILED_CONFIGURATION_ERROR

Severity	Description	Troubleshooting
major	HyperSwap reestablish failed. Either configuration of remote HyperSwap of peer '*local peer name*' on target '*target name*' does not match local configuration.	Make sure configuration on both machines is compatible and activate the HyperSwap. If problem persists contact IBM support.

HA_END_SYNC_FAILED_CONFIGURATION_ERROR

Severity	Description	Troubleshooting
major	Configuration of remote HyperSwap of peer '*local peer name*' on target '*target name*' does not match local configuration.	Make sure configuration on both machines is compatible and activate the HyperSwap. If problem persists contact IBM support.

HA_CHANGE_DESIGNATION

Severity	Description
informational	Local peer '*local peer name*' switched its designated role with peer on Target '*target name*'. It is now *designation*.

TARGET_QUORUM_WITNESS_NOT_LAPSED

Severity	Description
informational	Target '*target.name*' Quorum Witness '*quorum_witness.name*' is not lapsed.

TARGET_QUORUM_WITNESS_LAPSED

Severity	Description
Critical	Target '*target.name*' Quorum Witness '*quorum_witness.name*' is lapsed.

DATA_REBUILD_STARTED

Severity	Description
informational	Rebuild process started because system data is not protected. *data_percent*% of the data must be rebuilt.

DATA_REBUILD_COMPLETED

Severity	Description
informational	Rebuild process completed. System data is now protected.

DATA_REDIST_STARTED

Severity	Description
informational	Starting data transfer to new disks.

DATA_REDIST_COMPLETED

Severity	Description
informational	Completed data transfer to new disks.

DATA_REBUILD_COMPLETED_REDIST_STARTED

Severity	Description
informational	Rebuild process completed. System data is now protected. Starting data transfer to new disks.

DATA_REDIST_TIME_LIMIT_EXCEEDED

Severity	Description
warning	Data redistribution is taking too long. *data_percent*% of the required redistribution still remains to be done.

DATA_REDIST_BLOCKED

Severity	Description
informational	Blocking data transfer to new phased in media for *delay_in_seconds* seconds.

DATA_REDIST_ALLOWED

Severity	Description
informational	Allowing data transfer to new phased in media.

STORAGE_POOL_EXHAUSTED

Severity	Description	Troubleshooting
major	Pool '*pool*' is full. All volumes are locked.	Enlarge Storage Pool or move or delete volumes or Clones with Clone Deletion Priority 0.

STORAGE_POOL_UNLOCKED

Severity	Description
major	Pool '*pool*' has empty space. All volumes are unlocked.

STORAGE_POOL_SNAPSHOT_USAGE_INCREASED

Severity	Description
variable	Usage by snapshots of Storage Pool with name '*pool.name*' has reached *current%*.

STORAGE_POOL_SNAPSHOT_USAGE_DECREASED

Severity	Description
informational	Usage by snapshots of Storage Pool with name '*pool.name*' has decreased to *current%*.

HOST_CONNECTED

Severity	Description
informational	Host '*host*' has connected to the system.

HOST_DISCONNECTED

Severity	Description
warning	Host '*host*' has disconnected from the system.

HOST_MULTIPATH_OK

Severity	Description
informational	Host '*host*' has redundant connections to the system. #paths=*npaths*

HOST_NO_MULTIPATH_ONLY_ONE_PORT

Severity	Description
warning	Host '*host*' is connected to the system through only one of its ports. #paths=*npaths*

HOST_NO_MULTIPATH_ONLY_ONE_MODULE

Severity	Description
informational	Host '*host*' is connected to the system through only one interface module. #paths=*npaths*

SYSTEM_CAPACITY_USAGE_INCREASED

Severity	Description
variable	Usage of system physical capacity increased to *current*%.

SYSTEM_CAPACITY_USAGE_DECREASED

Severity	Description
informational	Usage of system physical capacity decreased to *current*%.

POOL_CREATE

Severity	Description
informational	Storage Pool of size *pool.size*GB*sparse_type* was created with name '*pool.name*'.

POOL_CREATE_FAILED_TOO_MANY

Severity	Description	Troubleshooting
warning	Storage Pool with name '*name*' could not be created. You are attempting to add more Storage Pools than the system permits.	Delete Storage Pools to allow new ones to be created.

POOL_RENAME

Severity	Description
informational	Storage Pool with name '*old_name*' was renamed '*pool.name*'.

POOL_RESIZE

Severity	Description
informational	Storage Pool with name '*pool.name*' was resized from size *old_size*GB*old_sparse_type* to *pool.size*GB*sparse_type*.

POOL_RESIZE_SNAPSHOTS

Severity	Description
informational	Snapshot size of Storage Pool with name '*pool.name*' was resized from size *old_size*GB to *pool.snapshot_size*GB.

POOL_CHANGE_LOCK_BEHAVIOR

Severity	Description
informational	Lock Behavior of Storage Pool with name '*pool.name*' is now '*state*'.

POOL_CONVERTED_TO_SPARSE

Severity	Description
informational	Storage Pool with name '*pool.name*' was converted to sparse.

POOL_CONVERTED_TO_REGULAR

Severity	Description
informational	Storage Pool with name '*pool.name*' was converted to regular.

POOL_CHANGE_PERF_CLASS

Severity	Description
informational	Performance Class of Storage Pool with name '*pool.name*' is now '*pool.perf_class*'.

POOL_CONFIG_SNAPSHOTS

Severity	Description
informational	Management policy of Mirroring snapshots of Storage Pool with name '*pool.name*' has changed'.

POOL_DELETE

Severity	Description
informational	Storage Pool with name '*pool.name*' was deleted.

FLASH_VDISK_TOO_SMALL

Severity	Description
major	Flash vdisk *name* on *Flash Enclosure* is too small.

FLASH_VDISK_LARGER_THAN_EXPECTED

Severity	Description
warning	Flash vdisk *name* on *Flash Enclosure* is larger than expected.

DATA_REDUCTION_TIER_IS_OFFLINE

Severity	Description
critical	Data reduction tier moved to offline mode

DATA_REDUCTION_RECOVERY_STARTED

Severity	Description
Informational	Data reduction recovery process started

DATA_REDUCTION_RECOVERY_FINISHED

Severity	Description
Informational	Data reduction recovery process ended

DATA_REDUCTION_RECOVERY_FAILED

Severity	Description
Minor	Data reduction recovery process failed with reason: '*reason*'

DATA_REDUCTION_RECOVERY_ABORT_STARTED

Severity	Description
Informational	Data reduction recovery abort started

DATA_REDUCTION_RECOVERY_ABORT_FINISHED

Severity	Description
Informational	Data reduction recovery abort finished

DATA_REDUCTION_TIER_IS_ONLINE

Severity	Description
informational	Data reduction tier moved to online state

DATA_REDUCTION_TIER_BECOMING_OFFLINE

Severity	Description
major	Data reduction tier is transitioning to offline mode

DATA_REDUCTION_RESUME_ONLINE_START

Severity	Description
informational	Data reduction is resuming from offline mode

DATA_REDUCTION_RESUME_ONLINE_FAILED

Severity	Description
warning	Data reduction failed resuming from offline mode

DATA_REDUCTION_COMPRESSION_ADAPTER_FAILED

Severity	Description
Major	Data Reduction node *reporting_node* reported about compression adapter failure and will be killed.

FLASH_COMPONENT_INITIALIZING

Severity	Description	Troubleshooting
informational	*Flash Component ID* initializing.	Contact IBM Support

FLASH_COMPONENT_OK

Severity	Description	Troubleshooting
informational	*Flash Component ID* status ok.	Contact IBM Support

FLASH_COMPONENT_FAILED

Severity	Description	Troubleshooting
variable	*Flash Component ID* has failed.	Contact IBM Support

FLASH_COMPONENT_REPLACED

Severity	Description	Troubleshooting
informational	*Component ID* was replaced. New serial *New Serial* Old serial *Old Serial*	Contact IBM Support

FLASH_FW_HOT_UPGRADE_STARTED

Severity	Description
informational	Start upgrade to version *version*. Enclosure id *component id*

FLASH_FW_HOT_UPGRADE_FINISHED

Severity	Description
informational	Finished upgrade to version *version*. Enclosure id *component id*

FLASH_UPGRADE_RESUMED

Severity	Description
informational	Finished upgrade resumed. Enclosure id *component id*

FLASH_FW_HOT_UPGRADE_RESUMED

Severity	Description
informational	Finished upgrade resumed. Enclosure id *component id*

FLASH_UPGRADE_STOPPED

Severity	Description
major	Flash upgrade stopped after *percents*%, reason: *reason*. Enclosure id *component id*

FLASH_FW_HOT_UPGRADE_STOPPED

Severity	Description
major	Flash upgrade stopped after *percents*%, reason: *reason*. Enclosure id *component id*

FLASH_FW_HOT_UPGRADE_FAILED

Severity	Description
major	Flash firmware hot upgrade failed, reason: *reason*. error: *ccl_error*. Enclosure id *component id*

FLASH_FW_HOT_UPGRADE_TIMEOUT

Severity	Description
major	Timeout while upgrading *component id*, progress: *percents*

FLASH_RAID_STATUS_CHANGED

Severity	Description
variable	Flash array raid status changed to *status*. Enclosure id *component id*.

FLASH_ARRAY_STATUS_CHANGED

Severity	Description
variable	Flash array status changed to *status*. Enclosure id *component id*.

FLASH_ENCRYPTION_STATUS_CHANGED

Severity	Description
informational	Encryption enabled changed to *encrypted*. Enclosure id *component id*.

FLASH_ENCRYPTION_ENABLE_NOT_ALLOWED

Severity	Description
major	Encryption enable not allowed: *reason*. Enclosure id *component id*.

FLASH_ENCRYPTION_DISABLE_NOT_ALLOWED

Severity	Description
major	Encryption disable not allowed: *reason*. Enclosure id *component id*.

ENCRYPT_ENABLE_FLASH_ENCLOSURE_FAILED

Severity	Description
major	Encryption enable failed for *component id*: *reason*

ENCRYPT_DISABLE_FLASH_ENCLOSURE_FAILED

Severity	Description
major	Encryption disable failed for *component id*: *reason*

FLASH_ENCRYPTION_UNLOCK_FAILED

Severity	Description
major	Flash Encryption unlock failed. Enclosure: *component id*.

FLASH_ENCLOSURE_WIPEOUT_FAILED

Severity	Description
major	Wipeout failed for *component id*: *reason*

FLASH_BBU_CHARGING_STATUS_CHANGED

Severity	Description
informational	BBU charging status changed to *status*. BBU: *component id*.

FLASH_BBU_CALIBRATION_STARTED

Severity	Description
informational	BBU calibration started, BBU: *component id*.

FLASH_BBU_CALIBRATION_STOPPED

Severity	Description
informational	BBU calibration stopped, BBU: *component id*.

FLASH_BBU_CALIBRATION_FAILED

Severity	Description
minor	BBU calibration failed, BBU: *component id*.

FLASH_CANISTER_CONNECTED_VIA_SERIAL_CABLE

Severity	Description
informational	Established serial connection with *component id*.

FLASH_CANISTER_CONNECTION_VIA_SERIAL_OK

Severity	Description
informational	Established serial connection with *component id*.

NO_CONNECTION_TO_FLASH_CANISTER_VIA_SERIAL

Severity	Description	Troubleshooting
warning	Failed to connect to *component id* via serial.	Contact IBM Support

FLASH_CANISTER_NO_CONNECTION_VIA_SERIAL

Severity	Description	Troubleshooting
warning	Failed to connect to *component id* via serial.	Contact IBM Support

NO_CONNECTION_TO_FLASH_CANISTER_VIA_ETH

Severity	Description	Troubleshooting
warning	Failed to connect to *component id* via ethernet.	Contact IBM Support

FLASH_CANISTER_NO_CONNECTION_VIA_ETH

Severity	Description	Troubleshooting
warning	Failed to connect to *component id* via ethernet.	Contact IBM Support

FLASH_CANISTER_CONNECTION_VIA_ETH_OK

Severity	Description	Troubleshooting
informational	Established connection to *component id* via ethernet.	Contact IBM Support

FLASH_CANISTER_IP_CHANGED

Severity	Description	Troubleshooting
informational	Changed connected ip of *component id* to *IP*.	Contact IBM Support

NO_CONNECTION_TO_FLASH_CANISTER_VIA_GW

Severity	Description	Troubleshooting
warning	Failed to connect to *component id* via *module id*.	Contact IBM Support

FLASH_CANISTER_NO_CONNECTION_VIA_GW

Severity	Description	Troubleshooting
warning	Failed to connect to *component id* via *module id*.	Contact IBM Support

FLASH_CANISTER_CONNECTED_VIA_GATEWAY

Severity	Description	Troubleshooting
informational	Established connection to *component id* via *module id*.	Contact IBM Support

FLASH_CANISTER_CONNECTION_VIA_GW_OK

Severity	Description	Troubleshooting
informational	Established connection to *component id* via *module id*.	Contact IBM Support

FLASH_CANISTER_ETH_LINK_MISWIRE

Severity	Description	Troubleshooting
minor	Flash canister *component id* is miswired.	Contact IBM Support

FLASH_CANISTER_ETH_LINK_WIRING_OK

Severity	Description	Troubleshooting
informational	Flash canister *component id* is no longer miswired.	Contact IBM Support

FLASH_CANISTER_GET_CONF_FAILED

Severity	Description	Troubleshooting
minor	Failed to get the configuration from *component id*.	Contact IBM Support

FLASH_CANISTER_GET_CONF_OK

Severity	Description	Troubleshooting
informational	Succeeded to get the configuration from *component id*.	Contact IBM Support

FLASH_CANISTER_IN_SERVICE_MODE

Severity	Description	Troubleshooting
major	Flash canister is in service mode *service*, canister: *component id*.	Contact IBM Support

FLASH_CANISTER_NO_LONGER_IN_SERVICE_MODE

Severity	Description
informational	Flash canister is no longer in service mode: *component id*.

FLASH_COMPONENT_TEMPERATURE_OK

Severity	Description
informational	Flash component *component id* temperature is within allowed limits.

FLASH_COMPONENT_TEMPERATURE_ABOVE_NORMAL

Severity	Description
warning	Flash component *component id* temperature is above normal.

FLASH_COMPONENT_TEMPERATURE_HIGH

Severity	Description
major	Flash component *component id* temperature is high, it exceeds operational level.

FLASH_COMPONENT_TEMPERATURE_CRITICALLY_HIGH

Severity	Description
critical	Flash component *component id* temperature is critical.

FLASH_ENCLOSURE_STATUS_CHANGED

Severity	Description	Troubleshooting
informational	*component id* status changed to *New Status*.	Contact IBM Support

FLASH_PSU_FAN_FAILED

Severity	Description
minor	*component id* fan has failed.

FLASH_PSU_FAN_OK

Severity	Description
informational	*component id* is ok.

FLASH_PSU_HAS_NO_INPUT_POWER

Severity	Description
major	PSU power supply has no input (AC) power. PSU id *component id*.

FLASH_PSU_HAS_INPUT_POWER

Severity	Description
informational	PSU power supply has input (AC) power. PSU id *component id*.

FLASH_ENCLOSURE_VERSION_IS_UNEXPECTED

Severity	Description
major	*component id* version is *version*, expected version is *expected*.

FLASH_ENCLOSURE_NEWER_VERSION_EXISTS

Severity	Description
minor	*component id* version *version* is supported. newer version *expected* exists.

FLASH_PSU_COMMUNICATION_ERROR

Severity	Description
minor	PSU communication error. PSU id *component id*.

FLASH_BBU_VPD_IS_NOT_VALID

Severity	Description
major	BBU VPD is not valid. BBU id *component id*.

FLASH_CARD_UNSUPPORTED

Severity	Description
warning	Flash card is unsupported. Flash card id *component id*.

FLASH_BBU_END_OF_LIFE

Severity	Description
warning	Battery is at end of life. BBU id *component id*.

FLASH_CARD_COMMUNICATION_ERROR

Severity	Description
minor	Flash card communication error. Flash card id *component id*.

FLASH_FAN_COMMUNICATION_ERROR

Severity	Description
minor	Fan communication error. Fan: *component id*.

FLASH_ENCLOSURE_THERMAL_THREASHOLD_EXCEEDED

Severity	Description
critical	Flash enclosure *component id* has shutdown after exceeding the thermal threshold.

FLASH_ENCLOSURE_ARRAY_OFFLINE

Severity	Description
critical	Flash enclosure *component id* array is offline.

FLASH_ENCLOSURE_STARTED_PHASEOUT

Severity	Description
informational	System started phasing out *Component ID*.

FLASH_ENCLOSURE_FINISHED_PHASEOUT

Severity	Description
informational	System finished phasing out *Component ID*.

FLASH_ENCLOSURE_STARTED_PHASEIN

Severity	Description
informational	System started phasing in *Component ID*.

FLASH_ENCLOSURE_FINISHED_PHASEIN

Severity	Description
informational	System finished phasing in *Component ID*.

FLASH_CR_KEY_SETUP_FAILED

Severity	Description
major	Failed to set challenge-response key on '*Component ID*'.

FLASH_CR_KEY_SETUP_OK

Severity	Description
informational	Challenge-response key was successfully set on '*Component ID*'.

FLASH_CR_KEY_SETUP_STARTED

Severity	Description
informational	Challenge-response key set started on '*Component ID*'.

FC_PORT_HAS_FAILED

Severity	Description	Troubleshooting
major	FC port service *port* has failed due to *code* (attempt number *Number of retries*)	Contact IBM Support

NTP_SERVER_TIME_DIFFERENCE_TOO_BIG

Severity	Description	Troubleshooting
warning	NTP server *NTP Server* sent a transaction with time difference of *Delta* seconds which exceeds the maximal difference of *Max Allowed* seconds. Transaction will be ignored, please check NTP server's and system's times.	Please contact your Administrator.

IPSEC_TUNNEL_OPENED

Severity	Description
informational	The IPSec tunnel named '*name*' between module *Module* and *Right IP* was opened

IPSEC_TUNNEL_CLOSED

Severity	Description
informational	The IPSec tunnel named '*name*' between module *Module* and *Right IP* was closed

IP_ACCESS_CANNOT_RESOLVE_ADDRESS

Severity	Description
informational	Cannot resolve address '*address*' added to the IP access group *IP access group name*.

IP_ACCESS_FAILED_SETTING_RULES

Severity	Description
informational	Failed setting IP access rules.

USB_ETHERNET_INTERFACE_OK

Severity	Description
informational	USB ethernet interface on module *Module* was reset successfully and is now OK.

USB_ETHERNET_INTERFACE_IS_STILL_DOWN

Severity	Description	Troubleshooting
minor	USB ethernet interface on module *Module* is down and failed to reset.	Contact IBM Support

USB_ETHERNET_INTERFACE_FAILED

Severity	Description	Troubleshooting
warning	USB ethernet interface on module *Module* failed.	Contact IBM Support

MIRROR_CREATE

Severity	Description
informational	A remote mirror was defined for Volume '*local volume name*'on Target '*target name*'. Remote Volume is '*remote volume name*'.

CG_MIRROR_CREATE

Severity	Description
informational	A remote mirror was defined for Consistency Group '*local CG name*'on Target '*target name*'. Remote Consistency Group is '*remote CG name*'.

MIRROR_CREATE_SLAVE

Severity	Description
informational	A remote mirror was defined by Target '*target name*' for Volume '*local volume name*'. Remote Volume is '*remote volume name*'.

CG_MIRROR_CREATE_SLAVE

Severity	Description
informational	A remote mirror was defined by Target '*target name*' for CG '*local CG name*'. Remote CG is '*remote CG name*'.

MIRROR_SCHEDULE_CHANGE

Severity	Description
informational	Schedule of remote mirror of '*local peer name*' is now '*schedule name*'.

MIRROR_CREATE_FAILED_TARGET_NOT_CONNECTED

Severity	Description	Troubleshooting
warning	Target could not be reached. Target with name '*target.name*' is currently not connected.	Connect the target system to this system.

REMOTE_OPERATION_FAILED_TIMED_OUT

Severity	Description	Troubleshooting
warning	Operation on remote machine timed out. Invoking '*Function Name*' on target '*Target Name*' timed out.	Retry operation. If problem persists contact IBM support.

MIRROR_RESYNC_FAILED

Severity	Description	Troubleshooting
major	Synchronization of meta data with mirror failed. Configuration of remote mirror of volume '*local volume name*' on target '*target name*' does not match local configuration.	Make sure configuration on both machines is compatible and activate the mirror. If problem persists contact IBM support.

MIRROR_RESYNC_FAILED_DUE_TO_THIN_PROVISIONING

Severity	Description	Troubleshooting
major	Synchronization of bitmaps with mirror failed. Not enough hard capacity left in Pool of volume '*mirror.local_volume_name*'.	Delete unnecessary volumes in pool or enlarge the pool's hard size.

MIRROR_SYNC_STARTED

Severity	Description
informational	Synchronization of remote mirror of volume '*local volume name*' on Target '*target name*' has started.

MIRROR_SYNC_ENDED

Severity	Description
informational	Synchronization of remote mirror of peer '*local peer name*' on target '*target name*' has ended.

MIRROR_CANNOT_CREATE_SYNC_JOB_ TOO_MANY_VOLUMES

Severity	Description
major	Synchronization of remote mirror of peer '*local peer name*' on target '*target name*' can not be synced , insufficent volume available for this operation.

MIRROR_CANNOT_CREATE_LRS_TOO_MANY_VOLUMES

Severity	Description
major	Synchronization of remote mirror of peer '*local peer name*' on target '*target name*' can not be synced , insufficent volume available for this operation.

MIRROR_REESTABLISH_FAILED_CONFIGURATION_ERROR

Severity	Description	Troubleshooting
major	Mirror reestablish failed. Either configuration of remote mirror of peer '*local peer name*' on target '*target name*' does not match local configuration.	Make sure configuration on both machines is compatible and activate the mirror. If problem persists contact IBM support.

MIRROR_ACTIVATE

Severity	Description
informational	The Remote Mirror of peer '*local peer name*' on Target '*target name*' was activated.

MIRROR_DEACTIVATE

Severity	Description
informational	The Remote Mirror of peer '*local peer name*' on Target '*target name*' was deactivated.

MIRROR_SLAVE_ACTIVATE

Severity	Description
informational	The mirror of peer '*local peer name*' on target '*target name*' was activated.

MIRROR_SLAVE_DEACTIVATE

Severity	Description
informational	The mirror of peer '*local peer name*' on target '*target name*' was deactivated.

MIRROR_DEACTIVATE_SECONDARY_LOCKED

Severity	Description
minor	The Remote Mirror of peer '*local peer name*' on Target '*target name*' was deactivated since the Pool on the secondary machine was locked.

MIRROR_DEACTIVATE_CONFIGURATION_ERROR

Severity	Description
minor	The Remote Mirror of peer '*local peer name*' on Target '*target name*' was deactivated since the Mirror configuration on the slave machine has changed.

MIRROR_DELETE

Severity	Description
informational	The Remote Mirror relation of peer '*local peer name*' to a peer on Target '*target name*' was deleted.

MIRROR_REVERSE_ROLE_TO_SLAVE

Severity	Description
informational	Local peer '*local peer name*' is now Slave of a peer on Target '*target name*'.

MIRROR_REVERSE_ROLE_TO_MASTER

Severity	Description
informational	Local peer '*local peer name*' is now Master of a peer on Target '*target name*'.

MIRROR_REVERSE_ROLE_OF_PEER_WITH_LCS_TO_MASTER

Severity	Description
informational	Local peer '*local peer name*' is now Master of a peer on Target '*target name*' external last consistent snapshot should be deleted manually .

MIRROR_SWITCH_ROLES_TO_SLAVE

Severity	Description
informational	Local peer '*local peer name*' switched roles with peer on Target '*target name*'. It is now Slave.

MIRROR_SWITCH_ROLES_TO_MASTER

Severity	Description
informational	Local peer '*local peer name*' switched roles with peer on Target '*target name*'. It is now Master.

MIRROR_REESTABLISH_FAILED_TOO_MANY_VOLUMES

Severity	Description	Troubleshooting
major	Last Consistent Snapshot of Slave peer '*local peer name*' could not be created. Maximal number of Volumes are already defined.	Delete Volumes to allow new ones to be created. Activate Mirror on the Master Machine.

MIRROR_END_SYNC_FAILED_CONFIGURATION_ERROR

Severity	Description	Troubleshooting
major	Configuration of remote mirror of peer '*local peer name*' on target '*target name*' does not match local configuration.	Make sure configuration on both machines is compatible and activate the mirror. If problem persists contact IBM support.

MIRROR_CHANGE_DESIGNATION

Severity	Description
informational	Local peer '*local peer name*' switched its designated role with peer on Target '*target name*'. It is now *designation*.

MIRROR_CANCEL_SNAPSHOT

Severity	Description
informational	All mirrored snapshots which were created for Mirror of peer '*local peer name*' and were not yet synchronized will not be mirrored in the remote machine.

DM_DEFINE

Severity	Description
informational	Data Migration was defined to Volume '*local volume name*' from Target '*target name*'.

DM_SYNC_STARTED

Severity	Description
informational	Migration to volume '*local volume name*' from Target '*target name*' has started.

DM_SYNC_ENDED

Severity	Description
informational	Migration to volume '*local volume name*' from target '*target name*' is complete.

DM_SYNC_ENDED_WITH_ERRORS

Severity	Description
Critical	Migration to volume '*local volume name*' from target '*target name*' has completed with *medium_errors_in_data_migration* error(s). Check previous events related to this volume for the list of affected LBAs.'.

DM_ACTIVATE

Severity	Description
informational	Migration to Volume '*local volume name*' from Target '*target name*' was activated.

DM_DEACTIVATE

Severity	Description
informational	Migration to Volume '*local volume name*' from Target '*target name*' was deactivated.

DM_DEACTIVATE_LUN_UNAVAILABLE

Severity	Description
minor	Migration to Volume '*local volume name*' from Target '*target name*' was deactivated since LUN is not available on one of the active paths to the target.

DM_DELETE

Severity	Description
informational	Definition of Data Migration to Volume '*local volume name*' from Target '*target name*' was deleted.

SCHEDULE_CREATE

Severity	Description
informational	Schedule was created with name '*schedule name*'.

SCHEDULE_UPDATE

Severity	Description
informational	Schedule with name '*schedule name*' was updated.

SCHEDULE_RENAME

Severity	Description
informational	Schedule with name '*old_name*' was renamed '*schedule name*'.

SCHEDULE_DELETE

Severity	Description
informational	Schedule with name '*schedule name*' was deleted.

MIRROR_RPO_OK

Severity	Description
informational	Mirror of local peer '*local peer name*' is now ahead of its specified RPO.

MIRROR_RPO_LAGGING

Severity	Description
informational	Mirror of local peer '*local peer name*' is now behind its specified RPO.

MIRROR_CHANGE_RPO

Severity	Description
informational	RPO or Mirror of local peer '*local peer name*' is now *RPO*.

MIRROR_IS_LAGGING_BEYOND_PERCENT_THRESHOLD

Severity	Description
warning	Last Replication Time of Mirror of local peer '*local peer name*' is *Last Replication Time*.

MIRROR_AUTO_FIX_REACHED_LIMIT

Severity	Description
warning	A remote checksum diff for mirror '*local peer name*' cannot be fixed automatically because we reached the auto fix limit.

MIRROR_IS_LAGGING_BEYOND_ABSOLUTE_THRESHOLD

Severity	Description
warning	Last Replication Time of Mirror of local peer '*local peer name*' is *Last Replication Time*.

MIRROR_INCOMPATIBLE_VERSION_FOR_UNMAP_SUPPORT

Severity	Description
warning	Mirror of peer '*local peer name*' on target '*target name*' cannot support unmap, remote machine has incompatible version.

XMIRROR_DEFINE

Severity	Description
informational	A xmirror master '*xmirror name*' was defined for volume '*local volume name*'.

XMIRROR_DEFINE_SLAVE

Severity	Description
informational	A xmirror slave '*xmirror name*' was defined for volume '*local volume name*'.

XMIRROR_DEFINE_SMASTER

Severity	Description
informational	A xmirror smaster '*xmirror name*' was defined for volume '*local volume name*'.

XMIRROR_ACTIVATE

Severity	Description
informational	Xmirror '*xmirror name*' was activated.

XMIRROR_DEACTIVATE

Severity	Description
informational	Xmirror '*xmirror name*' was deactivated.

XMIRROR_DELETE

Severity	Description
informational	Xmirror '*xmirror name*' was deleted.

XMIRROR_CHANGE_SLAVE_ROLE_TO_MASTER

Severity	Description
informational	Xmirror '*xmirror name*' was changed to standalone.

XMIRROR_CHANGE_MASTER_ROLE_TO_SLAVE

Severity	Description
informational	Xmirror '*xmirror name*' was changed to slave.

XMIRROR_CHANGE_SMASTER_ROLE_TO_MASTER

Severity	Description
informational	Xmirror '*xmirror name*' was changed to master.

XMIRROR_CHANGE_MASTER_ROLE_TO_SMASTER

Severity	Description
informational	Xmirror '*xmirror name*' was changed to smaster.

XMIRROR_STANDBY_MIRROR_REGISTERED

Severity	Description
informational	Xmirror '*xmirror name*' registered a standby mirror on SMASTER system

XMIRROR_COMPROMISED

Severity	Description
warning	Xmirror '*xmirror name*' is compromised, reason: *Compromise Reason*

XMIRROR_RESTORED

Severity	Description
informational	Xmirror '*xmirror name*' restored after being compromised

XMIRROR_RENAMED

Severity	Description
informational	Xmirror '*Old Xmirror Name*' was renamed to '*xmirror name*'.

MAP_VOLUME

Severity	Description
informational	Volume with name '*volume.name*' was mapped to LUN '*LUN*' for *host_or_cluster* with name '*host*'.

MAP_PROXY_VOLUME

Severity	Description
informational	IBM Hyper-Scale Mobility Volume with name '*name*' was mapped to LUN '*LUN*' for *host_or_cluster* with name '*host*'.

UNMAP_VOLUME

Severity	Description
informational	Volume with name '*volume.name*' was unmapped from *host_or_cluster* with name '*host*'.

UNMAP_PROXY_VOLUME

Severity	Description
informational	IBM Hyper-Scale Mobility Volume with name '*name*' was unmapped from *host_or_cluster* with name '*host*'.

SPECIAL_TYPE_SET

Severity	Description
informational	Type of *host_or_cluster* with name '*host*' was set to '*type*'.

SERVICE_HAS_FAILED

Severity	Description	Troubleshooting
variable	*Component ID* has failed.	Contact IBM Support

SERVICE_FAILED_TO_PHASEIN

Severity	Description	Troubleshooting
major	*Component ID* failed to phase-in.	Contact IBM Support

SERVICE_FAILED_TO_RESTART

Severity	Description	Troubleshooting
major	*Component ID* failed to restart.	Contact IBM Support

MODULE_FAILED

Severity	Description	Troubleshooting
critical	*Component ID* failed.	Contact IBM Support

DISK_HAS_FAILED

Severity	Description	Troubleshooting
variable	Disk *Component ID* Failed.	Please contact your Administrator.

SSD_HAS_FAILED

Severity	Description	Troubleshooting
major	SSD *Component ID* Failed.	Please contact your Administrator.

VAULT_DEVICE_HAS_FAILED

Severity	Description	Troubleshooting
minor	Vault device *Component ID* Failed.	Please contact your Administrator.

COMPONENT_TEST_OF_DISK_HAS_FAILED

Severity	Description	Troubleshooting
major	Test of *Component ID* has failed with error *Error*.	Please contact your Administrator.

COMPONENT_TEST_OF_SSD_HAS_FAILED

Severity	Description	Troubleshooting
major	Test of *Component ID* has failed with error *Error*.	Please contact your Administrator.

COMPONENT_TEST_OF_BOOT_MEDIA_HAS_FAILED

Severity	Description	Troubleshooting
major	Test of *Component ID* has failed with error *Error*.	Please contact your Administrator.

COMPONENT_TEST_OF_VAULT_DEVICE_HAS_FAILED

Severity	Description	Troubleshooting
major	Test of *Component ID* has failed with error *Error*.	Please contact your Administrator.

BOOT_MEDIA_COMPONENT_TEST_FAILED

Severity	Description	Troubleshooting
major	Test of *Component ID* has failed with error *Error*.	Please contact your Administrator.

DISK_STARTED_PHASEOUT

Severity	Description
informational	System started phasing out *Component ID*.

DISK_STARTED_PHASEIN

Severity	Description
informational	System started phasing in *Component ID*.

DISK_FINISHED_PHASEIN

Severity	Description
informational	System finished phasing in *Component ID*.

DISK_FINISHED_PHASEOUT

Severity	Description
informational	System finished phasing out *Component ID*.

DISK_RECOVERED

Severity	Description
critical	Disk *Component ID* is functioning again.

MODULE_STARTED_PHASEOUT

Severity	Description
informational	System started phasing out *Component ID*.

MODULE_FINISHED_PHASEOUT

Severity	Description
informational	System finished phasing out *Component ID*.

MODULE_STOPPED_PHASEOUT_ DUE_TO_MANAGEMENT_REQUIREMENT

Severity	Description
major	System stopped phasing out *Component ID* due to management requirement.

START_WORK

Severity	Description
informational	System has entered ON state.

SYSTEM_HAS_ENTERED_MAINTENANCE_MODE

Severity	Description
warning	System has entered MAINTENANCE state [*Reason*]

SYSTEM_LEFT_CHARGING_STATE

Severity	Description
informational	System is sufficiently charged.

USER_SHUTDOWN

Severity	Description
major	System is shutting down due to a user request.

EMERGENCY_SHUTDOWN_NOW

Severity	Description	Troubleshooting
critical	System is shutting down in emergency shutdown mode due to: *Emergency Shutdown Reason*.	Please contact your Administrator.

SHOULD_HAVE_BEEN_EMERGENCY_SHUTDOWN

Severity	Description
critical	An emergency shutdown has been detected, but emergency shutdown is disabled for the detected reason. Shutdown reason: *Shutdown Reason*.

DATA_SERVICE_STARTED_PHASEOUT

Severity	Description
informational	System started phasing out *Component ID*.

DATA_SERVICE_STARTED_PHASEIN

Severity	Description
informational	System started phasing in *Component ID*.

DATA_SERVICE_FINISHED_PHASEIN

Severity	Description
informational	System finished phasing in *Component ID*.

DATA_SERVICE_FINISHED_PHASEOUT

Severity	Description
informational	System finished phasing out *Component ID*.

TXN_REBUILD_STARTED

Severity	Description
informational	Start rebuild process of txns.

TXN_REBUILD_ENDED

Severity	Description
informational	End rebuild process for txns.

TXN_REDIST_STARTED

Severity	Description
informational	Start redist process for txns.

TXN_REDIST_ENDED

Severity	Description
informational	End redist process for txns.

TXN_REMOVE_NODE_WAS_EXECUTED

Severity	Description
informational	Dev command txn_remove_node was called for node *Node ID*.

TXN_REMOVE_SLICE_IMAGE_WAS_EXECUTED

Severity	Description
informational	Dev command txn_remove_slice_image was called for slice *Slice*, position *Position*.

DISK_MARKED_TO_PHASEOUT

Severity	Description
informational	System started phasing out *Component ID*.

DISK_MARKED_TO_PHASEIN

Severity	Description
informational	System started phasing out *Component ID*.

CANNOT_CREATE_NEW_DATA_DISTRIBUTION

Severity	Description
informational	System cannot phaseout disks for storage medium *Storage Medium*.

CANNOT_RESIZE_FLASH_MEDIUM_POOLS

Severity	Description
informational	System failed to resize flash medium pools for TMS phaseout.

DATA_PROTECTION_STATUS_CHANGED

Severity	Description
variable	Data protection status has changed from '*old_status*' to '*new_status*'

VAULT_DEVICE_SECURE_ERASE_PROCESS_SUCCESSFUL

Severity	Description
informational	Vault device secure erase process successful.

VAULT_DEVICE_SECURE_ERASE_PROCESS_FAILED

Severity	Description
major	Vault device secure erase process failed [*Reason*]

VAULT_DEVICE_FAILED_SECURE_ERASE

Severity	Description
major	Secure erase for *Component ID* failed. [*Reason*].

SYSTEM_PHYSICAL_CAPACITY_CHANGED

Severity	Description
informational	System physical capacity is now *Capacity*GB.

SYSTEM_EFFECTIVE_CAPACITY_CHANGED

Severity	Description
informational	System effective capacity is now *Capacity*GB.

SYSTEM_OUT_OF_PHYSICAL_SPACE

Severity	Description
critical	System has run out of physical capacity. All volumes are now write-protected.

SYSTEM_NORMAL_OPERATION_RESUMED

Severity	Description
informational	Normal operation is resumed. Volumes have been restored to their original write-protection state.

ENCRYPT_ENABLE_STARTED

Severity	Description
informational	Starting encryption activation. This process can take several minutes to complete.

ENCRYPT_ENABLE_COMPLETED

Severity	Description
informational	Encryption is in effect.

ENCRYPT_ENABLE_NOT_COMPLETED

Severity	Description	Troubleshooting
major	Cannot complete encryption activation because of *reason*. *Count* vault device(s) and *Count* flash enclosure(s) could not be enrolled.	Please contact technical support

ENCRYPT_DISABLE_STARTED

Severity	Description
informational	Starting encryption deactivation. This process can take several minutes to complete.

ENCRYPT_DISABLE_COMPLETED

Severity	Description
informational	Encryption is no longer in effect.

ENCRYPT_DISABLE_NOT_COMPLETED

Severity	Description	Troubleshooting
major	Cannot complete encryption deactivation because of *reason*. *Count* vault device(s) and *Count* flash enclosure(s) could not be crypto erased.	Please contact technical support

ENCRYPT_KEYSERVER_ADDED

Severity	Description
informational	A key server named '*Key Server Name*' was added.

ENCRYPT_KEYSERVER_DELETED

Severity	Description
informational	Key server '*Key Server Name*' was deleted.

ENCRYPT_KEYSERVER_EDITED

Severity	Description
informational	Details of key server '*Key Server Name*' were modified.

ENCRYPT_KEYSERVER_RENAMED

Severity	Description
informational	Key server '*Old Name*' was renamed to '*New Name*'.

ENCRYPT_KEYSERVER_CHECK_STATUS_STARTED

Severity	Description
informational	Start checking connectivity status of the keyservers currently defined in the system. This process can take several minutes to complete.

ENCRYPT_KEYSERVER_CHECK_STATUS_COMPLETED

Severity	Description
informational	Completed checking connectivity status of the keyservers currently defined in the systemt.

ENCRYPT_KEYSERVER_REKEY_COMPLETED

Severity	Description
informational	Key server '*Key Server Name*' rekey completed.

ENCRYPT_LOCAL_REKEY_COMPLETED

Severity	Description
informational	Local key rekey completed.

ENCRYPT_CHANGE_KEY_SCHEME_COMPLETED

Severity	Description
informational	Change key scheme from external to local key completed.

ENCRYPT_CHANGE_KEY_SCHEME_FAILED

Severity	Description
major	Change key scheme from external to local key failed because *failure reason*.

ENCRYPT_CHANGE_KEY_SCHEME_ROLLBACK_FAILED

Severity	Description
major	Cannot rollback change key scheme.

ENCRYPT_UNABLE_TO_UPDATE_KEY_ DURING_DEACTIVATE_ON_KEYSERVER

Severity	Description
major	Could not update key server '*Keyserver Name*' regarding encryption deactivation. Please check key server status.

ENCRYPT_KEYSERVER_REKEY_FAILED

Severity	Description
major	Cannot complete rekey with key server '*Key Server Name*'.

ENCRYPT_LOCAL_REKEY_FAILED

Severity	Description
major	Local key rekey failed because *failure reason*.

ENCRYPT_KEYSERVER_REKEY_ROLLBACK_FAILED

Severity	Description
major	Cannot rollback failed rekey with key server '*Key Server Name*'.

ENCRYPT_LOCAL_REKEY_ROLLBACK_FAILED

Severity	Description
major	Cannot rollback failed local key rekey.

ENCRYPT_RECOVERY_KEY_ENTERED

Severity	Description
informational	Valid recovery key share was entered by user '*User Name*'.

ENCRYPT_INVALID_RECOVERY_KEY_ENTERED

Severity	Description
major	Invalid recovery key share was entered by user '*User Name*'.

ENCRYPT_RECOVERY_KEYS_GENERATED

Severity	Description
informational	Recovery keys created.

ENCRYPT_RECOVERY_KEY_REKEY_SUCCESS

Severity	Description
informational	Recovery key rekey was successful.

ENCRYPT_RECOVERY_KEY_REKEY_FAIL

Severity	Description
major	Recovery key rekey failed.

ENCRYPT_RECOVERY_KEY_VERIFIED

Severity	Description
informational	Recovery key verifed successfully for user '*User Name*'.

ENCRYPT_RECOVERY_KEY_VERIFY_FAILED

Severity	Description
major	Recovery key verification failed for user '*User Name*'.

ENCRYPT_RECOVERY_KEY_ALL_SHARES_VERIFIED

Severity	Description
informational	All recovery key shares have been verified.

ENCRYPT_KR_WRITE_FAILED

Severity	Description	Troubleshooting
critical	Key repository write failed with error code: *rc*.	Please contact technical support

ENCRYPT_KR_READ_FAILED

Severity	Description	Troubleshooting
major	Key repository read failed with error code: *rc*.	Please contact technical support

ENCRYPT_UNABLE_TO_RETRIEVE_KEY_FROM_KEYSERVER

Severity	Description
major	Failed to retrieve key from key server '*Keyserver Name*' via *TEXT* on module *node id*. Please verify that the key server type and version are supported. If so, please check its status.

ENCRYPT_RECOVERY_KEY_RECOVER_SUCCESSFUL

Severity	Description
informational	Key recovery was successful, unlocking system.

ENCRYPTION_CERTIFICATE_FOR_XIV_IS_NOT_INSTALLED

Severity	Description	Troubleshooting
critical	XIV certificate is not installed.	Check output of pki_list for a certificate named XIV and contact technical support

ENCRYPT_UNABLE_TO_DELETE_MASTER_KEYSERVER

Severity	Description
informational	Deletion of master key server '*Keyserver Name*' is not allowed. Please define another key server as master first'.

ENCRYPTION_SKMIP_ERROR

Severity	Description	Troubleshooting
major	Module *Module* reported *Keyserver Name* returned error:*error code - TEXT*	Please contact the next level of support.

WIPEOUT_STARTED

Severity	Description
informational	Starting wipeout process. This process can take several minutes to complete.

WIPEOUT_COMPLETED

Severity	Description
informational	Wipeout process finished successfully.

WIPEOUT_NOT_COMPLETED

Severity	Description	Troubleshooting
major	Cannot complete wipeout process because of *reason*. *Count* vault device(s) and *Count* flash enclosure(s) could not be crypto erased.	Please contact technical support

DIMM_FAILED

Severity	Description	Troubleshooting
major	*Component ID* has failed. Hardware status: *Status*.	Contact IBM Support

CPU_FAILED

Severity	Description	Troubleshooting
major	*Component ID* has failed. Hardware status: *Status*.	Contact IBM Support

NIC_FAILED

Severity	Description	Troubleshooting
major	*Component ID* has failed. Hardware status: *Status*.	Contact IBM Support

MODULE_BBU_FAILED

Severity	Description	Troubleshooting
major	*BBU id* has failed. Hardware status: '*Status*'. BBU state: '*State*'.	Contact IBM Support

MODULE_BBU_OK

Severity	Description
informational	*BBU id* is now OK.

DIMM_WAS_REMOVED

Severity	Description	Troubleshooting
minor	The DIMM with serial number '*Serial*' was removed from *ModuleId*.	Was this DIMM actually removed?

PSU_CHANGE_DETECTED

Severity	Description	Troubleshooting
informational	*Component ID* has been changed from a serial number '*old_serial*', part number '*old_part_number*', to serial number '*new_serial*' and part number '*new_part_number*'.	Was this PSU actually replaced?

PSU_WAS_REMOVED

Severity	Description	Troubleshooting
warning	*Component ID* with a serial number '*Serial*' and part number '*Part Number*' was removed from the system.	Was this PSU actually removed?

PSU_MISSING_FROM_INSTALL

Severity	Description	Troubleshooting
major	System was installed without *Component ID*.	This PSU was not in the system at install time.

PSU_WAS_INSTALLED

Severity	Description	Troubleshooting
informational	*Component ID* with a serial number '*Serial*' and part number '*Part Number*' was installed in the system.	Was this PSU actually installed?

NIC_CHANGE_DETECTED

Severity	Description	Troubleshooting
major	*Component ID* has been changed from a serial of *old_serial* to *new_serial*.	Was this NIC actually replaced?

MODULE_BBU_IS_TOO_OLD

Severity	Description
major	*BBU id* installed more than *max use time* months ago on *insert time*.

MODULE_BBU_SHELF_TIME_EXCEEDED

Severity	Description
major	*BBU id* installed more than *max shelf time* months after manufacturing date *manufacturing date*.

TECHNICIAN_WORK_STARTED

Severity	Description
informational	Technician work has started, expected to end at *End Time*. Comment: *Comment*.

TECHNICIAN_WORK_ENDED

Severity	Description
informational	Technician work has ended after *Elapsed Time* minutes. Comment: *Comment*.

TECHNICIAN_WORK_TIMED_OUT

Severity	Description
warning	Technician work has timed out after *Elapsed Time* minutes. Comment: *Comment*.

XIV_SUPPORT_ENABLED

Severity	Description
informational	XIV support access from *From* is enabled from *Start Time* until *Finish Time*. Comment: *Comment*.

XIV_SUPPORT_ENABLED_NO_TIME_LIMIT

Severity	Description
informational	XIV support access from *From* is enabled from *Start Time* until explicitly disabled. Comment: *Comment*.

XIV_SUPPORT_DISABLED

Severity	Description
informational	XIV support access is disabled.

XIV_SUPPORT_WINDOW_TIMEOUT

Severity	Description
informational	XIV support work window timeout is expired.

HOST_DEFINE

Severity	Description
informational	Host of type *host.type* was defined with name '*host.name*'.

HOST_UPDATE

Severity	Description
informational	Host named '*host.name*' was updated.

CLUSTER_CREATE

Severity	Description
informational	Cluster was defined with name '*cluster.name*'.

HOST_DEFINE_FAILED_TOO_MANY

Severity	Description	Troubleshooting
warning	Host with name '*name*' could not be defined. You are attempting to define more hosts than the system permits.	Delete Hosts to allow new ones to be defined.

CLUSTER_CREATE_FAILED_TOO_MANY

Severity	Description	Troubleshooting
warning	Cluster with name '*name*' could not be defined. You are attempting to define more Clusters than the system permits.	Delete Clusters to allow new ones to be defined.

HOST_RENAME

Severity	Description
informational	Host with name '*old_name*' was renamed '*host.name*'.

CLUSTER_RENAME

Severity	Description
informational	Cluster with name '*old_name*' was renamed '*cluster.name*'.

HOST_DELETE

Severity	Description
informational	Host with name '*host.name*' was deleted.

CLUSTER_DELETE

Severity	Description
informational	Cluster with name '*cluster.name*' was deleted.

HOST_ADD_PORT

Severity	Description
informational	Port of type *type* and ID '*port_name*' was added to Host with name '*host.name*'.

CLUSTER_ADD_HOST

Severity	Description
informational	Host with name '*host.name*' was added to Cluster with name '*cluster.name*'.

HOST_REMOVE_PORT

Severity	Description
informational	Port of type *type* and ID '*port_name*' was removed from Host with name '*host.name*' was deleted.

CLUSTER_REMOVE_HOST

Severity	Description
informational	Host with name '*host.name*' was removed from Cluster with name '*cluster.name*'.

DESTINATION_DEFINE

Severity	Description
informational	Destination with name '*name*' was defined.

DESTINATION_UPDATE

Severity	Description
informational	Destination with name '*name*' was updated.

DESTINATION_DELETE

Severity	Description
informational	Destination with name '*name*' was deleted.

DESTINATION_RENAME

Severity	Description
informational	Destination with name '*old name*' was renamed '*new name*'.

DESTINATION_GROUP_CREATE

Severity	Description
informational	Destination Group with name '*name*' was created.

DESTINATION_GROUP_UPDATE

Severity	Description
informational	Destination Group with name '*name*' was updated.

DESTINATION_GROUP_DELETE

Severity	Description
informational	Destination Group with name '*name*' was deleted.

DESTINATION_GROUP_RENAME

Severity	Description
informational	Destination Group with name '*old name*' was renamed '*new name*'.

DESTINATION_GROUP_ADD_DESTINATION

Severity	Description
informational	Destination with name '*destination name*' was added to destination group '*destgroup name*'.

DESTINATION_GROUP_REMOVE_DESTINATION

Severity	Description
informational	Destination with name '*destination name*' was removed from destination group '*destgroup name*'.

RULE_CREATE

Severity	Description
informational	Rule with name '*name*' was created.

RULE_UPDATE

Severity	Description
informational	Rule with name '*name*' was updated.

RULE_DELETE

Severity	Description
informational	Rule with name '*name*' was deleted.

RULE_RENAME

Severity	Description
informational	Rule with name '*old name*' was renamed '*new name*'.

SMTP_GATEWAY_DEFINE

Severity	Description
informational	SMTP gateway with name '*name*' was defined.

SMTP_GATEWAY_UPDATE

Severity	Description
informational	SMTP gateway with name '*name*' was updated.

SMTP_GATEWAY_DELETE

Severity	Description
informational	SMTP gateway with name '*name*' was deleted.

SMTP_GATEWAY_RENAME

Severity	Description
informational	SMTP gateway with name '*old name*' was renamed '*new name*'.

SMTP_GATEWAY_PRIORITIZE

Severity	Description
informational	SMTP gateways were prioritized; the new order is *order*.

CALL_HOME_CONNECTION_OK

Severity	Description
informational	Events are sent to the Call Home server by SMTP gateway '*name*'.

CALL_HOME_CONNECTION_PROBLEM

Severity	Description	Troubleshooting
major	Events are not sent to the Call Home server by SMTP gateway '*name*'. Reason: *Event Reason*.	Please contact IBM support.

SMTP_GATEWAY_FAILED

Severity	Description
major	SMTP gateway with name '*name*' has failed. It will not be used until *Retry Time*.

SMTP_GATEWAY_VIA_NODE_FAILED

Severity	Description
warning	Sending event *Event Code* (*Event Index*) through *SMTP Gateway* via *Module ID* has failed; Error message: '*Error Message*'.

SMS_GATEWAY_DEFINE

Severity	Description
informational	SMS gateway with name '*name*' was defined.

SMS_GATEWAY_UPDATE

Severity	Description
informational	SMS gateway with name '*name*' was updated.

SMS_GATEWAY_DELETE

Severity	Description
informational	SMS gateway with name '*name*' was deleted.

SMS_GATEWAY_RENAME

Severity	Description
informational	SMS gateway with name '*old name*' was renamed '*new name*'.

SMS_GATEWAY_PRIORITIZE

Severity	Description
informational	SMS gateways were prioritized; the new order is *order*.

CONS_GROUP_CREATE

Severity	Description
informational	Consistency Group with name '*cg.name*' was created.

CONS_GROUP_CREATE_FAILED_TOO_MANY

Severity	Description	Troubleshooting
warning	Consistency Group with name '*name*' could not be created. You are attempting to add more Consistency Groups than the system permits.	Delete Consistency Groups to allow new ones to be created.

CONS_GROUP_RENAME

Severity	Description
informational	Consistency Group with name '*old_name*' was renamed '*cg.name*'.

SECONDARY_CONS_GROUP_RENAME

Severity	Description
informational	Consistency Group with name '*old_name*' was renamed '*cg.name*' by primary machine.

CONS_GROUP_DELETE

Severity	Description
informational	Consistency Group with name '*cg.name*' was deleted.

CONS_GROUP_ADD_VOLUME

Severity	Description
informational	Volume with name '*volume.name*' was added to Consistency Group with name '*cg.name*'.

SLAVE_CONS_GROUP_ADD_VOLUME

Severity	Description
informational	Volume with name '*volume.name*' was added to Consistency Group with name '*cg.name*' by its remote peer.

CONS_GROUP_REMOVE_VOLUME

Severity	Description
informational	Volume with name '*volume.name*' was removed from Consistency Group with name '*cg.name*'.

SLAVE_CONS_GROUP_REMOVE_VOLUME

Severity	Description
informational	Volume with name '*volume.name*' was removed from Consistency Group with name '*cg.name*' by its remote peer.

CONS_GROUP_SNAPSHOTS_CREATE

Severity	Description
informational	Snapshot Group for Consistency Group with name '*cg.name*' was created with name '*sg.name*'.

CONS_GROUP_SNAPSHOTS_CREATE_FAILED_TOO_MANY

Severity	Description	Troubleshooting
warning	Snapshot Group for Consistency Group '*cg.name*' could not be created. You are attempting to add more Volumes than the system permits.	Delete Volumes to allow new ones to be created.

CONS_GROUP_SNAPSHOTS_OVERWRITE

Severity	Description
informational	Snapshot Group named '*sg.name*' was overriden for Consistency Group with name '*cg.name*'.

SLAVE_CONS_GROUP_SNAPSHOTS_CREATE

Severity	Description
informational	Mirrored Snapshot Group for Consistency Group with name '*cg.name*' was created with name '*sg.name*'.

SLAVE_CONS_GROUP_SNAPSHOTS_OVERWRITE

Severity	Description
informational	Mirrored Snapshot Group named '*sg.name*' was overriden for Consistency Group with name '*cg.name*'.

MIRROR_CONS_GROUP_SNAPSHOTS_CREATE

Severity	Description
informational	Mirrored Snapshot Group for Consistency Group with name '*cg.name*' was created with name '*sg.name*'.

MIRROR_CONS_GROUP_SNAPSHOTS_OVERWRITE

Severity	Description
informational	Mirrored Snapshot Group named '*sg.name*' was overriden for Consistency Group with name '*cg.name*'.

REMOTE_MIRROR_CONS_GROUP_SNAPSHOTS_ NOT_CREATED_YET

Severity	Description
minor	Remote Mirrored Snapshot Group for Consistency Group with name '*cg.name*' on Target '*target name*' were not created yet.

MIRROR_SNAPGROUP_CREATE_FAILED

Severity	Description
minor	Remote snapshot group named '*snapshot group name*' was not created successfully. Error code is '*error*'

SNAPSHOT_GROUP_RESTORE

Severity	Description
informational	Volumes were restored from Snapshot Group with name '*sg.name*'.

SNAPSHOT_GROUP_RENAME

Severity	Description
informational	Snapshot Group with name '*old_sg.name*' were renamed to '*sg.name*'.

SNAPSHOT_GROUP_DUPLICATE

Severity	Description
informational	All Snapshots in Snapshot Group with name '*sg.name*' were duplicated. Duplicate Snapshot Group is named '*sg.name*'.

SNAPSHOT_GROUP_FORMAT

Severity	Description
informational	All Snapshots in Snapshot Group with name '*sg.name*' were formatted'.

SNAPSHOT_GROUP_DELETE

Severity	Description
informational	All Snapshots in Snapshot Group with name '*sg.name*' were deleted.

SNAPSHOT_GROUP_CHANGE_PRIORITY

Severity	Description
informational	Deletion Priority of all Snapshots in Snapshot Group with name '*sg.name*' were changed from '*old priority*' to '*new priority*'.

SNAPSHOT_GROUP_LOCK

Severity	Description
informational	All Snapshots in Snapshot Group with name '*sg.name*' were locked.

SNAPSHOT_GROUP_UNLOCK

Severity	Description
informational	All Snapshots in Snapshot Group with name '*sg.name*' were unlocked.

SNAPSHOT_GROUP_DELETED_DUE_TO_POOL_EXHAUSTION

Severity	Description
informational	All Snapshots in Snapshot Group with name '*snapshot.sg_name*' have been deleted because Storage Pool with name '*snapshot.pool_name*' is full.

SNAPSHOT_GROUP_DISBAND

Severity	Description
informational	Snapshot Group with name '*sg.name*' was dismantled. All Snapshots which belonged to that Snapshot Group should be accessed directly.

CONS_GROUP_MOVE

Severity	Description
informational	Consistency Group with name '*cg.name*' has been moved from Storage Pool '*orig_pool.name*' to Pool '*pool.name*'.

XCG_CREATE

Severity	Description
informational	Cross Consistency Group with name '*xcg*' was created.

XCG_DELETE

Severity	Description
informational	Cross Consistency Group with name '*xcg*' was deleted.

XCG_ADD_CG

Severity	Description
informational	CG with name '*cg.name*' was added to Cross Consistency Group with name '*xcg*'.

XCG_REMOVE_CG

Severity	Description
informational	CG with name '*cg.name*' was removed from Cross Consistency Group with name '*xcg*'.

TARGET_DEFINE

Severity	Description
informational	Target was defined named '*target.name*'.

TARGET_DEFINE_FAILED_TOO_MANY

Severity	Description	Troubleshooting
warning	Target could not be defined. You are attempting to define more targets than the system permits.	Delete targets to allow new ones to be defined.

TARGET_RENAME

Severity	Description
informational	Target named '*old_name*' was renamed '*target.name*'.

TARGET_DELETE

Severity	Description
informational	Target named '*target.name*' was deleted.

TARGET_ALLOW_ACCESS

Severity	Description
informational	Target '*target.name*' is allowed to access this machine.

TARGET_PORT_ADD

Severity	Description
informational	Port '*port_name*' was added to target named '*target.name*'.

TARGET_PORT_REMOVE

Severity	Description
informational	Port '*port_name*' was removed from target named '*target.name*'.

TARGET_PORT_ACTIVATE

Severity	Description
informational	Port '*port_name*' in target named '*target.name*' was activated.

TARGET_PORT_DEACTIVATE

Severity	Description
informational	Port '*port_name*' was deactivated in target named '*target.name*'.

TARGET_CONNECTIVITY_CREATE

Severity	Description
informational	Port '*Connection Remote Port Address*' of target named '*Connection Target Name*' is connected to the system through *Local FC Port*.

TARGET_ISCSI_CONNECTIVITY_CREATE

Severity	Description
informational	Port '*Connection Remote Port Address*' of target named '*Connection Target Name* is connected to the system through ip interface '*Local IP interface*'.

TARGET_CONNECTIVITY_CREATE_FAILED_TOO_MANY

Severity	Description	Troubleshooting
warning	Port could not be connected to the system. You are attempting to define more connections than the system permits.	Delete Connections to allow new ones to be created.

TARGET_CONNECTIVITY_DELETE

Severity	Description
informational	Port '*Connection Remote Port Address*' of target named '*Connection Target Name*' was disconnected from *Local FC Port*.

TARGET_ISCSI_CONNECTIVITY_DELETE

Severity	Description
informational	Port '*Connection Remote Port Address*' of target named '*Connection Target Name*' was disconnected from ip interface '*Local IP interface*'.

TARGET_CONNECTIVITY_ACTIVATE

Severity	Description
informational	Connectivity between Port '*Connection Remote Port Address*' of target named '*Connection Target Name*' and *Local FC Port* was activated.

TARGET_ISCSI_CONNECTIVITY_ACTIVATE

Severity	Description
informational	Connectivity between Port '*Connection Remote Port Address*' of target named '*Connection Target Name*' and ip interface '*Local IP interface*' was activated.

TARGET_CONNECTIVITY_DEACTIVATE

Severity	Description
informational	Connectivity between Port '*Connection Remote Port Address*' of target named '*Connection Target Name*' and *Local FC Port* was deactivated.

TARGET_ISCSI_CONNECTIVITY_DEACTIVATE

Severity	Description
informational	Connectivity between Port '*Connection Remote Port Address*' of target named '*Connection Target Name*' and ip interface '*Local IP interface*' was deactivated.

TARGET_CONNECTIVITY_CONFLICT_DETECTED

Severity	Description
major	Connectivity between Port '*Connection Remote Port Address*' of target named '*Connection Target Name*' and FC port '*Local IP interface*' will be deleted due to a connectivity conflict.

TARGET_ISCSI_CONNECTIVITY_CONFLICT_DETECTED

Severity	Description
major	Connectivity between Port '*Connection Remote Port Address*' of target named '*Connection Target Name*' and IP interface '*Local IP interface*' will be deleted due to a connectivity conflict.

TARGET_CONNECTION_ESTABLISHED

Severity	Description
informational	Target named '*target.name*' is accessible through remote service *module_id*.

TARGET_CONNECTION_DISCONNECTED

Severity	Description
minor	Target named '*target.name*' is no longer accessible through remote service *module_id*.

TARGET_DISCONNECTED

Severity	Description
major	Target named '*target.name*' is no longer accessible through any gateway module.

TARGET_CLOCK_SKEW_ABOVE_LIMIT

Severity	Description
warning	Target '*target.name*' has clock skew above the allowed limit relative to local machine.

TARGET_CLOCK_SKEW_RESOLVED

Severity	Description
informational	Target named '*target.name*' clock skew has been resolved.

TARGET_LINK_DOWN_BEYOND_THRESHOLD

Severity	Description
major	Target named '*target.name*' is not accessible for a long time.

OLVM_DELETE_ALL_REFERENCES_TO_SOURCE

Severity	Description
major	Target named '*target.name*' was released from all IBM Hyper-Scale Mobility relationships.

TARGET_SYNC_RATE_CHANGED

Severity	Description
informational	Target '*target.name*' sync rate changed. max_initialization_rate: '*target.max_initialization_rate*', max_resync_rate: '*target.max_resync_rate*', max_syncjob_rate: '*target.max_syncjob_rate*'.

TARGET_ADD_QUORUM_WITNESS

Severity	Description
Informational	Target '*target_name*' added quorum witness '*quorum_witness_name*'.

TARGET_REMOVE_QUORUM_WITNESS

Severity	Description
Informational	Target '*target_name*' removed quorum witness '*quorum_witness_name*'.

TARGET_SYSTEM_DETAILS_UPDATED

Severity	Description
Informational	Target named '*target_name*' has updated details: old system id '*old_system_id*', old machine serial '*old_machine_serial*', new system id '*new_system_id*', new machine serial '*new_machine_serial*'.

TARGET_HANDSHAKE_COMPLETED

Severity	Description
Informational	Target named '*target_name*' has completed handshake.

TARGET_HANDSHAKE_FAILED

Severity	Description
Major	Target named '*target_name*' handshake failed with reason '*Failure Reason*'

TARGET_HANDSHAKE_REINITIATED

Severity	Description
Informational	Target named '*target_name*' reinitiated handshake process.

SNAPSHOT_CREATE

Severity	Description
informational	Snapshot named '*snapshot.name*' was created for volume named '*volume.name*'.

SNAPSHOT_DELETE

Severity	Description
informational	Snapshot with name '*snapshot.name*' was deleted.

SNAPSHOT_OVERWRITE

Severity	Description
informational	Snapshot named '*snapshot.name*' was overriden for volume named '*volume.name*'.

SNAPSHOT_FORMAT

Severity	Description
informational	Snapshot named '*snapshot.name*' was formatted.

SNAPSHOT_CREATE_FAILED_TOO_MANY

Severity	Description	Troubleshooting
warning	Snapshot for volume named '*volume.name*' could not be created. You are attempting to add more volumes than the system permits.	Delete volumes to allow new ones to be created.

SNAPSHOT_DUPLICATE

Severity	Description
informational	Snapshot named '*snapshot.name*' was created as duplicate of Snapshot named '*original_snapshot.name*'.

SNAPSHOT_DUPLICATE_FAILED_TOO_MANY

Severity	Description	Troubleshooting
warning	Snapshot named '*snapshot.name*' could not be duplicated. You are attempting to add more volumes than the system permits.	Delete volumes to allow new ones to be created.

SNAPSHOT_RESTORE

Severity	Description
informational	Volume named '*volume.name*' was restored from Snapshot named '*snapshot.name*'.

SNAPSHOT_CHANGE_PRIORITY

Severity	Description
informational	Snapshot Delete Priority of Snapshot named '*snapshot.name*' was changed from '*old_priority*' to '*snapshot.delete_priority*'.

SNAPSHOT_DELETED_DUE_TO_POOL_EXHAUSTION

Severity	Description
warning	Snapshot named '*snap.name*' has been deleted because Storage Pool named '*snap.pool_name*' is full.

MIRROR_SNAPSHOT_CREATE

Severity	Description
informational	Mirrored Snapshot named '*snapshot.name*' was created for volume named '*volume.name*'.

MIRROR_SNAPSHOT_CREATE_FAILED

Severity	Description
minor	Remote snapshot named '*snapshot name*' was not created successfully. Error code is '*error*'

MIRROR_SNAPSHOT_OVERWRITE

Severity	Description
informational	Mirrored Snapshot named '*snapshot.name*' was overriden for volume named '*volume.name*'.

MIRROR_SLAVE_SNAPSHOT_CREATE

Severity	Description
informational	Mirrored Snapshot named '*snapshot.name*' was created for volume named '*volume.name*'.

MIRROR_SLAVE_SNAPSHOT_OVERWRITE

Severity	Description
informational	Mirrored Snapshot named '*snapshot.name*' was overriden for volume named '*volume.name*'.

MEDIUM_ERROR_IN_DATA_MIGRATION

Severity	Description	Troubleshooting
critical	Medium error in data migration into volume '*Volume Name*' at LBA *LBA* for *Length* blocks.	Remote machine indicated Medium Error when read.

USER_DEFINED

Severity	Description
informational	A user with name '*Name*' and category *Category* was defined.

USER_DELETED

Severity	Description
informational	A user with name '*Name*' and category *Category* was deleted.

USER_RENAMED

Severity	Description
informational	User with name '*Old Name*' was renamed '*New Name*'.

USER_UPDATED

Severity	Description
informational	User with name '*Name*' was updated.

USER_ADDED_TO_USER_GROUP

Severity	Description
informational	User '*User Name*' was added to user group '*User Group Name*'.

USER_REMOVED_FROM_USER_GROUP

Severity	Description
informational	User '*User Name*' was removed from user group '*User Group Name*'.

USER_GROUP_CREATED

Severity	Description
informational	A user group with name '*Name*' was created.

USER_GROUP_DELETED

Severity	Description
informational	A user group with name '*Name*' was deleted.

USER_GROUP_RENAMED

Severity	Description
informational	User group with name '*Old Name*' was renamed '*New Name*'.

LDAP_AUTHENTICATION_ACTIVATED

Severity	Description
informational	LDAP authentication activated.

LDAP_AUTHENTICATION_DEACTIVATED

Severity	Description
warning	LDAP authentication deactivated.

LDAP_CONFIGURATION_CHANGED

Severity	Description
warning	LDAP configuration has changed.

LDAP_CONFIGURATION_RESET

Severity	Description
warning	LDAP configuration has reset.

USER_LOGIN_HAS_SUCCEEDED

Severity	Description
informational	User '*User Name*' from IP '*Client Address*' successfully logged into the system.

USER_LOGIN_HAS_FAILED

Severity	Description
warning	User '*User Name*' from IP '*Client Address*' failed logging into the system.

USER_HAS_FAILED_TO_RUN_COMMAND

Severity	Description
warning	User '*User Name*' from IP '*Client Address*' failed authentication when trying to run command '*Command Line*'.

LDAP_SERVER_INACCESSIBLE

Severity	Description
minor	LDAP server *FQDN* is inaccessible.

LDAP_SERVER_ACCESSIBLE

Severity	Description
informational	LDAP server *FQDN* is now accessible.

LDAP_SSL_CERTIFICATE_ABOUT_TO_EXPIRE

Severity	Description
variable	SSL Certificate of LDAP server '*Server FQDN*' is about to expire on *Expiration Date* (*Counter* notification).

LDAP_SERVER_WAS_ADDED

Severity	Description
informational	LDAP server '*Server FQDN*' was added to the system.

LDAP_SERVER_WAS_REMOVED

Severity	Description
informational	LDAP server '*Server FQDN*' was removed from the system.

DESIGNATED_MSM_USER

Severity	Description
informational	*Description*

DOMAIN_POLICY_SET

Severity	Description
informational	Domain policy for *Parameter Name* set to '*Parameter Value*'

USER_ADDED_TO_DOMAIN

Severity	Description
informational	User *User Name* was added to domain *Domain Name* (*Exclusive*).

USER_REMOVED_FROM_DOMAIN

Severity	Description
informational	User *User Name* was removed from domain *Domain Name*.

APPADMIN_CAPABILITIES_SET

Severity	Description
informational	Application admin capabilities have been set to *Capabilities*

ACCESS_TO_HOST_GRANTED_TO_USER_GROUP

Severity	Description
informational	User group '*User Group Name*' was granted access to host '*Host Name*'.

ACCESS_OF_USER_GROUP_TO_HOST_REMOVED

Severity	Description
informational	Access of User group '*User Group Name*' to host '*Host Name*' was removed.

ACCESS_TO_CLUSTER_GRANTED_TO_USER_GROUP

Severity	Description
informational	User group '*User Group Name*' was granted access to cluster '*Cluster Name*'.

ACCESS_OF_USER_GROUP_TO_CLUSTER_REMOVED

Severity	Description
informational	Access of User group '*User Group Name*' to cluster '*Cluster Name*' was removed.

COMPONENT_TEST_HAS_FAILED

Severity	Description	Troubleshooting
variable	Test of *Component ID* has failed. Failure reason: *Failure Reason*.	Contact IBM Support

COMPONENT_TEST_SUCCEEDED

Severity	Description	Troubleshooting
informational	Test of *Component ID* succeeded.	Contact IBM Support

MODULE_COMPONENT_TEST_STARTED

Severity	Description
informational	Test of *Component ID* started.

DISK_COMPONENT_TEST_STARTED

Severity	Description
informational	Test of *Component ID* started.

IB_SWITCH_COMPONENT_TEST_STARTED

Severity	Description
informational	Test of *Component ID* started.

SSD_COMPONENT_TEST_STARTED

Severity	Description
informational	Test of *Component ID* started.

VAULT_DEVICE_COMPONENT_TEST_STARTED

Severity	Description
informational	Test of *Component ID* started.

BOOT_MEDIA_COMPONENT_TEST_STARTED

Severity	Description
informational	Test of *Component ID* started.

FLASH_COMPONENT_TEST_STARTED

Severity	Description
informational	Test of *Component ID* started.

BOOT_MEDIA_FAILED

Severity	Description	Troubleshooting
major	*Component ID* has failed. Hardware status: *Status*.	Contact IBM Support

MODULE_BBU_TEST_STARTED

Severity	Description
informational	Test of *Component ID* started.

MODULE_BBU_TEST_WILL_CONTINUE

Severity	Description
warning	Test of *Component ID* will continue. After waiting *Minutes* minutes, capacity is still *Capacity*%.

COMPONENT_WAS_PHASED_OUT

Severity	Description
informational	*Component ID* was phased-out.

COMPONENT_WAS_FAILED

Severity	Description
variable	Component *Component ID* was marked as failed.

COMPONENT_FAILURE_WAS_CANCELED

Severity	Description
informational	Component *Component ID* failure status was reset.

COMPONENT_WAS_PHASED_IN

Severity	Description
informational	*Component ID* was phased-in.

COMPONENT_WAS_EQUIPPED

Severity	Description
informational	*Component ID* was equipped.

INTERFACE_SERVICES_ACTIVATED

Severity	Description
informational	Interface services of *Module ID* were activated.

COMPONENT_FIRMWARE_UPGRADE_ABORTING

Severity	Description
warning	Aborting *Upgrade type* upgrade of *Firmware type* firmware, version *Label*, on *Scope*. Abort reason: *Reason*. Waiting for current upgrade item to complete.

COMPONENT_FIRMWARE_UPGRADE_ABORTED

Severity	Description
warning	Aborted *Upgrade type* upgrade of *Firmware type* firmware, version *Label*, on *Scope*. Abort reason: *Reason*. Progress *Attempted/Total*, *Successes* succeeded, *Failures* failed, *No-Ops* no-ops.

COMPONENT_FIRMWARE_UPGRADE_DONE

Severity	Description
informational	Finished *Upgrade type* upgrade of *Firmware type* firmware, version *Label*, on *Scope*. *Successes* succeeded, *Failures* failed, *No-Ops* no-ops.

COMPONENT_FIRMWARE_UPGRADE_STARTED

Severity	Description
informational	Starting *Upgrade type* upgrade of *Firmware type* firmware, version *Label*, on *Scope*.

COMPONENT_FIRMWARE_CANNOT_PHASEOUT_COMPONENT

Severity	Description
minor	Cannot phase out *Component ID*: *Error*. Firmware upgrade result was: *Upgrade result*.

COMPONENT_FIRMWARE_CANNOT_FAIL_COMPONENT

Severity	Description
minor	Cannot fail *Component ID*: *Error*. Firmware upgrade result was: *Upgrade result*.

MIRRORING_CONNECTIVITY_TO_NON_XIV_TARGET

Severity	Description	Troubleshooting
warning	Gateway Node #*Node ID*: connection to *target name*:*target's connection index* mirroring connection was established, but being ignored because the remote end is not an XIV target or is not properly configured	Please make sure the target's designation is correct, that the connection's parameters identify the intended system and that the intended system has a target_port defined for this system.

DM_CONNECTIVITY_TO_XIV_TARGET

Severity	Description	Troubleshooting
warning	Gateway Node #*Node ID*: connection to *target name*:*target's connection index* DM connection was established, but being ignored because the remote end is an XIV target configured for mirroring, rather than a host	Please make sure the target's designation is correct, that the connection's parameters identify the intended system and that the intended system has a host defined for this system (and not a target_port).

EMERGENCY_ROOT_ACCESS

Severity	Description
warning	Emergency login to 'root' account on module '*Component ID*' from '*IP Address*' using key number '*Authorized Key Number*'.

EMERGENCY_CONSOLE_ACCESS

Severity	Description
warning	Emergency login to '*Unix Account Name*' account on module '*Component ID*' from tty '*TTY Device*'.

CR_BYPASS_ACCESS

Severity	Description
warning	*Command that bypasses CR mechanism* access to '*Unix Account Name*' account on module '*Component ID*' from '*IP Address*'.

CR_KEY_SETUP_OK

Severity	Description
informational	Challenge-response key was successfully set on all modules in the system.

CR_KEY_UPGRADE_NOT_DONE

Severity	Description
warning	Challenge-response key was not upgraded on the system since a valid key has been previously set.

CR_KEY_SETUP_FAILED

Severity	Description
major	Failed to set challenge-response key on module '*Component ID*'.

SSH_REVOKE_KEY_OK

Severity	Description
informational	Authorized SSH key ending with '*Tail of Authorized SSH key*' was successfully revoked for user '*Unix Account Name*' on all modules in the system.

SSH_REVOKE_KEY_FAILED

Severity	Description
major	Failed to revoke authorized SSH key ending with '*Tail of Authorized SSH key*' for user '*Unix Account Name*' on module '*Component ID*'.

IB_SWITCH_PHASEOUT_STARTED

Severity	Description
informational	System started phasing out *Component ID*.

IB_SWITCH_PHASEIN_STARTED

Severity	Description
informational	System started phasing in *Component ID*.

IB_SWITCH_PHASEIN_FAILED

Severity	Description
warning	*Component ID* has failed to phase-in.

IB_SWITCH_CONFIG_FAILED

Severity	Description	Troubleshooting
warning	*Component ID* could not be configured	Contact IBM Support

IB_SWITCH_FIRMWARE_INCOMPATIBLE

Severity	Description	Troubleshooting
warning	The firmware version of *Component ID* is '*New Version*'. It should be '*Old Version*'.	None

IB_SWITCH_CPLD_INCOMPATIBLE

Severity	Description	Troubleshooting
warning	The CPLD version of *Component ID* is invalid.	None

IB_SWITCH_FIRMWARE_UPDATE_IN_PROGRESS

Severity	Description	Troubleshooting
informational	Firmware version of *Component ID* is '*Old Version*'. It should be '*New Version*'. Firmware will be updated. It may take a while.	Wait for IB switch to complete initialization.

IB_SWITCH_FIRMWARE_UPDATED

Severity	Description	Troubleshooting
informational	The firmware version of *Component ID* was updated to '*New Version*'.	None.

IB_SWITCH_LOG_COLLECT_OK

Severity	Description
informational	Log collection for IB switch '*switch_id*' completed successfuly. Log can be found in module '*log_module*' in the following directory: '*log_location*'.

IB_SWITCH_LOG_COLLECT_FAILED

Severity	Description
warning	Log collection for IB switch '*switch_id*' failed. Failure reason: '*failure_reason*'.

IB_SWITCH_MGMT_LINK_AVAILABLE

Severity	Description
informational	Management link *Type* of *Component ID* is available.

IB_SWITCH_MGMT_LINK_UNAVAIL

Severity	Description
warning	Management link *Type* of *Component ID* is unavailable.

IB_SWITCH_MGMT_LINK_MISWIRED

Severity	Description	Troubleshooting
warning	Management link *Type* of *Component ID* is connected to *Router* instead of *Expected Router*.	Check wiring

IB_SWITCH_MGMT_AVAILABLE

Severity	Description
informational	Management of *Component ID* is available.

IB_SWITCH_MGMT_UNAVAIL

Severity	Description
major	Management of *Component ID* is unavailable.

IB_SWITCH_PSU_OK

Severity	Description
informational	*Component ID* has returned to normal state.

IB_SWITCH_PSU_MONITOR_FAILED

Severity	Description
minor	*Component ID* sensor cannot be read.

IB_SWITCH_PSU_IS_MISSING

Severity	Description
minor	*Component ID* is not present.

IB_SWITCH_PSU_FAIL

Severity	Description
major	*Component ID* failed.

IB_SWITCH_BBU_OK

Severity	Description
informational	*Component ID* has returned to normal state.

IB_SWITCH_BBU_MONITOR_FAILED

Severity	Description
minor	*Component ID* sensor cannot be read.

IB_SWITCH_BBU_IS_MISSING

Severity	Description
minor	*Component ID* is not present.

IB_SWITCH_BBU_FAIL

Severity	Description
major	*Component ID* failed.

IB_SWITCH_FAN_OK

Severity	Description
informational	*Component ID* has returned to normal state.

IB_SWITCH_FAN_MONITOR_FAILED

Severity	Description
minor	*Component ID* sensor cannot be read.

IB_SWITCH_FAN_IS_MISSING

Severity	Description
minor	*Component ID* is not present.

IB_SWITCH_FAN_FAIL

Severity	Description
major	*Component ID* failed.

IB_SWITCH_PSU_FAN_OK

Severity	Description
informational	*Component ID* has returned to normal state.

IB_SWITCH_PSU_FAN_MONITOR_FAILED

Severity	Description
minor	*Component ID* sensor cannot be read.

IB_SWITCH_PSU_FAN_IS_MISSING

Severity	Description
minor	*Component ID* is not present.

IB_SWITCH_PSU_FAN_FAIL

Severity	Description
major	*Component ID* failed.

IB_SWITCH_VOLTAGE_MONITOR_FAILED

Severity	Description
minor	*Component ID* voltage sensor *Sensor Type* cannot be read.

IB_SWITCH_VOLTAGE_CHANGE

Severity	Description
informational	*Component ID* voltage sensor *Sensor Type* changed from *Old Status* to *Status*.

IB_SWITCH_PSU_VOLTAGE_MONITOR_FAILED

Severity	Description
minor	*Component ID* voltage sensor cannot be read.

IB_SWITCH_PSU_VOLTAGE_CHANGE

Severity	Description
informational	*Component ID* voltage sensor changed from *Old Status* to *Status*.

IB_SWITCH_BBU_VOLTAGE_MONITOR_FAILED

Severity	Description
minor	*Component ID* voltage sensor cannot be read.

IB_SWITCH_BBU_VOLTAGE_CHANGE

Severity	Description
informational	*Component ID* voltage sensor changed from *Old Status* to *Status*.

IB_SWITCH_TEMPERATURE_MONITOR_FAILED

Severity	Description
minor	*Component ID* sensor *Sensor Type* temperature cannot be read.

IB_SWITCH_TEMPERATURE_OK

Severity	Description
informational	*Component ID* sensor *Sensor Type* temperature has returned to normal state.

IB_SWITCH_TEMPERATURE_HIGH

Severity	Description
minor	*Component ID* sensor *Sensor Type* temperature is high.

IB_SWITCH_TEMPERATURE_CRITICAL

Severity	Description
major	*Component ID* sensor *Sensor Type* temperature is critical.

IB_SWITCH_PSU_TEMPERATURE_MONITOR_FAILED

Severity	Description
minor	*Component ID* temperature cannot be read.

IB_SWITCH_PSU_TEMPERATURE_OK

Severity	Description
informational	*Component ID* temperature has returned to normal state.

IB_SWITCH_PSU_TEMPERATURE_HIGH

Severity	Description
minor	*Component ID* temperature is high.

IB_SWITCH_PSU_TEMPERATURE_CRITICAL

Severity	Description
major	*Component ID* temperature is critical.

IB_SWITCH_BBU_TEMPERATURE_MONITOR_FAILED

Severity	Description
minor	*Component ID* temperature cannot be read.

IB_SWITCH_BBU_TEMPERATURE_OK

Severity	Description
informational	*Component ID* temperature has returned to normal state.

IB_SWITCH_BBU_TEMPERATURE_HIGH

Severity	Description
minor	*Component ID* temperature is high.

IB_SWITCH_BBU_TEMPERATURE_CRITICAL

Severity	Description
major	*Component ID* temperature is critical.

IB_PORT_MOVED

Severity	Description
informational	Infinibind module port '*module_port*' moved from '*from_port*' to '*to_port*'.

SYSTEM_TEMPERATURE_IS_ABOVE_NORMAL

Severity	Description	Troubleshooting
warning	System temperature is *System Temperature*C, which is above the normal temperature.	Cool the system down.

SYSTEM_TEMPERATURE_IS_HIGH

Severity	Description	Troubleshooting
minor	System temperature is *System Temperature*C, which is high.	Cool the system down.

SYSTEM_TEMPERATURE_IS_CRITICALLY_HIGH

Severity	Description	Troubleshooting
critical	System temperature is *System Temperature*C, which exceeds operational level. Please initiate shutdown sequence. Without further action, the system will automatically shut itself down if it reaches *Shutdown Threshold*C.	Cool the system down immediately or shut down the system using 'shutdown -y' and contact IBM support.

SYSTEM_TEMPERATURE_IS_CRITICALLY_HIGH_ SHUTTING_DOWN

Severity	Description	Troubleshooting
critical	System temperature is *System Temperature*C, which is critically high. Shutting down the system.	Shut down the system using 'shutdown -y' and contact IBM support.

SYSTEM_TEMPERATURE_IS_TOO_HIGH

Severity	Description	Troubleshooting
major	System temperature is *System Temperature*C. It approaches the maximal allowable value.	Cool the system down and contact IBM support.

SYSTEM_TEMPERATURE_IS_TOO_LOW

Severity	Description	Troubleshooting
major	System temperature is *System Temperature*C, which is lower than the minimal allowable value.	Contact IBM Support

SYSTEM_TEMPERATURE_IS_OK_NOW

Severity	Description
informational	System temperature is *System Temperature*C, which is within allowed limits.

SYSTEM_AVERAGE_POWER_PREPARATION_STARTED

Severity	Description
informational	System average power consumption preparation has started, when it's over the up-to-date value will be available.

SYSTEM_AVERAGE_POWER_PREPARATION_OVER

Severity	Description
informational	System average power consumption preparation is over, you can now read the up-to-date value.

SYSTEM_AVERAGE_POWER_PREPARATION_FAILED

Severity	Description
warning	System average power consumption preparation has failed, try again later.

ENCRYPT_ENABLE_DRIVE_FAILED

Severity	Description	Troubleshooting
major	Failed to enable encryption for *Component ID*. Error code: *Failure Reason*.	Contact IBM Support

ENCRYPT_ENABLE_VAULT_DEVICE_FAILED

Severity	Description	Troubleshooting
major	Failed to enable encryption for *Component ID*. Error code: *Failure Reason*.	Contact IBM Support

VAULT_DEVICE_ENCRYPTING_ENABLE_FAILED

Severity	Description	Troubleshooting
major	Failed to enable encryption for *Component ID*. Error code: *Failure Reason*.	Contact IBM Support

VAULT_DEVICE_SECURE_ERASE_FAILED

Severity	Description	Troubleshooting
major	Failed to secure erase *Component ID*. Error code: *Failure Reason*.	Contact IBM Support

VAULT_DEVICE_SECURE_ERASE_SUCCESSFUL

Severity	Description	Troubleshooting
informational	Secure erase was successful for *Component ID*.	Contact IBM Support

VAULT_DEVICE_SECURE_ERASE_NOT_DONE

Severity	Description	Troubleshooting
informational	Secure erase was not done for *Component ID* as it is unsecured.	Contact IBM Support

MODULE_BBU_OVERHEATING

Severity	Description	Troubleshooting
major	*bbu* reached a temperature of *temp*C, above *critical_set*C. BBU will be disconnected until it cools down below *critical_release*C.	Contact IBM Support

MODULE_BBU_IS_DISCHARGING

Severity	Description
informational	*BBU id* changed state from '*old_state*' to '*new state*'.

MODULE_BBU_STOPPED_DISCHARGING

Severity	Description
informational	*BBU id* changed state from '*old_state*' to '*new state*'.

MODULE_BBU_CHARGING_WAS_EXPLICITLY_ENABLED

Severity	Description
minor	*BBU id* was not charging, it had to be reset explicitly.

MODULE_BBU_NOT_CHARGING_AFTER_RESET

Severity	Description
major	*BBU id* is still not charging after *Reset Attempts* reset attempts.

MODULE_BBU_NOT_CHARGED_AFTER_RESET

Severity	Description
major	*BBU id* was reset and it is charging, but after *Minutes* minutes it is still only *Percent Charged*% charged, which is not enough.

MODULE_BBU_STILL_NOT_SUFFICIENTLY_CHARGED

Severity	Description
major	*BBU id* is charging, but after *Minutes* minutes it is still only *Percent Charged*% charged, which is not enough.

MODULE_BBU_DISCHARGING_WAS_EXPLICITLY_ENABLED

Severity	Description
minor	*BBU id* was in discharge disabled mode, it had to be enabled explicitly.

MODULE_BBU_IS_CHARGING

Severity	Description
informational	*BBU id* changed state from '*old_state*' to '*new state*'.

MODULE_BBU_IS_FULL

Severity	Description
informational	*BBU id* changed state from '*old_state*' to '*new state*'.

MODULE_BBU_DRIVER_NOT_LOADED

Severity	Description	Troubleshooting
major	BBU driver is not loaded in module *Component*. modprobe ruby	Contact IBM Support

MODULE_BBU_CONTROLLER_NOT_PRESENT

Severity	Description	Troubleshooting
major	BBU controller board not detected in module *Component*.	Contact IBM Support

MODULE_IS_MISSING_EPOW_CABLE

Severity	Description	Troubleshooting
major	The EPOW cable in the BBU controller board on module *Component* is not detected.	Contact IBM Support

MODULE_EPOW_CABLE_OK_NOW

Severity	Description
informational	The EPOW cable in the BBU controller board on module *Component* is now OK.

MODULE_IS_MISSING_POWER_SENSE_CABLE

Severity	Description	Troubleshooting
major	The power sense cable in the BBU controller board on module *Component* is not detected.	Contact IBM Support

MODULE_POWER_SENSE_CABLE_OK_NOW

Severity	Description
informational	The power sense cable in the BBU controller board on module *Component* is now OK.

MODULE_BBU_CALIBRATION_STARTED

Severity	Description
informational	*BBU id* started calibration.

MODULE_BBU_CALIBRATION_ENDED

Severity	Description
informational	*BBU id* ended calibration with status '*result*'.

MODULE_BBU_TEMPERATURE_TOO_HIGH_ FOR_CALIBRATION

Severity	Description	Troubleshooting
major	*BBU id* temperature is *Temperature.Temperature Tenths*C which is too high. Calibration is stopped.	Contact IBM Support

MODULE_BBU_TEST_IN_CHARGING_MODE

Severity	Description
informational	Test of *Component ID* is pending and will resume once it gets to a capacity of *Target Capacity*%. The current capacity is *Current Capacity*%.

SDR_PSU_STATUS_OK

Severity	Description
informational	*Psu* (location *Location*) is now OK. Changed from '*previous_sdr_status*' to '*sdr_status*'.

SDR_PSU_STATUS_BAD

Severity	Description
warning	*PSU* (location *Location*) is failed or off. Changed from '*previous_sdr_status*' to '*current_sdr_status*'.

INVALID_PSU_PART_NUMBER

Severity	Description	Troubleshooting
major	*PSU* has an invalid part number '*PN*'.	Please contact IBM support and have the PSU replaced.

CMOS_BATTERY_TOO_WEAK

Severity	Description	Troubleshooting
major	The CMOS battery on *Module* is too weak.	Please contact IBM support and have the battery replaced.

CMOS_BATTERY_IS_OK

Severity	Description
informational	The CMOS battery on *Module* is now OK.

FC_LINK_IS_NOW_DOWN

Severity	Description	Troubleshooting
major	FC port *Component* Active Firmware *Firmware version* - link disconnected.	Contact IBM Support

FC_LINK_IS_NOW_UP

Severity	Description	Troubleshooting
informational	FC port *Component* - link regained.	Contact IBM Support

FC_LINK_SYNC_ERROR

Severity	Description	Troubleshooting
major	FC port *Component* - errors on the physical layer: *Reason*.	Please contact support.

FC_PORT_TEST_STARTED

Severity	Description
informational	FC port *Component* - test started

FC_PORT_TEST_NOT_STARTED

Severity	Description	Troubleshooting
informational	FC port *Component* - test not started.	Check port state

FC_PORT_TEST_FAILED

Severity	Description	Troubleshooting
major	FC port *Component* - test failed.	Contact IBM Support

FC_PORT_TEST_SUCCESS

Severity	Description
informational	FC port *Component* - test success.

FC_PORT_TEST_ABORTED

Severity	Description
informational	FC port *Component* - test aborted.

COMPONENT_NETWORK_LINK_IS_DOWN

Severity	Description	Troubleshooting
major	Network interface to *Connected Component* on *Component ID* - link disconnected.	Contact IBM Support

COMPONENT_NETWORK_LINK_IS_UP

Severity	Description	Troubleshooting
informational	Network interface to component *Connected Component* on *Component ID* - link regained.	Contact IBM Support

MODULE_IS_MISSING_REQUIRED_MEMORY

Severity	Description	Troubleshooting
major	*Component ID* has less memory (*actual_mem* GB) than is defined for use (*req_mem* GB).	Please contact your Administrator.

POD_IB_MISWIRE

Severity	Description	Troubleshooting
warning	POD module miswired: *Module ID*.	Contact IBM Support

POD_IB_MISWIRE_CORRECTED

Severity	Description
informational	POD module miswire corrected: *Module ID*.

IMM_USB_INTERFACE_FAILED

Severity	Description	Troubleshooting
minor	IMM USB interface on module *Module* failed and can't be reset.	Contact IBM Support

MODULE_SET_LED_LOCATOR_FAILED

Severity	Description
warning	Failed to set LED locator on *module*.

MODULE_SET_LED_LOCATOR_COMPLETED

Severity	Description
informational	LED locator set successfully on *module*.

PERF_CLASS_RESOURCE_EXHAUSTION

Severity	Description
warning	Exhausted all allowed resources for performance classes on *Module Id*, BUSY until resources available.

CONNECTED_HOSTS_LIMIT_REACHED

Severity	Description
informational	Number of connected Hosts was reached for port '*port_id*' in Module *Module Id*.

QoS_HAS_BEEN_TRIGGERED

Severity	Description
informational	Queues on port '*port_id*' in Module *Module Id* caused QoS to be activated.

PERF_CLASS_RATE_AT_LIMIT

Severity	Description
informational	Performance class '*perf_class*' on *Module Id* reached its limit of *Limit Limit Name*, IOs being throttled.

INDEPENDENT_PERF_CLASS_RATE_AT_LIMIT

Severity	Description
informational	Performance class '*perf_class*' object *type*:*name* on *Module Id* reached its limit of *Limit Limit Name*, IOs being throttled.

PORT_PREP_FOR_UPGRADE_TIMED_OUT

Severity	Description
warning	Preparation of *port_type* port '*local_port_name*' for hot-upgrade timed out due to host '*host_name*' port '*host_port_name*'*host_port_addr*

INTERFACE_DISCONNECTED_FROM_TARGET

Severity	Description
major	Interface node on module *module* cannot access target '*target*' through any gateway module.

INTERFACE_RECONNECTED_TO_TARGET

Severity	Description
major	Interface node on module *module* can access target '*target*'.

METADATA_SERVICE_DB_CREATE

Severity	Description
informational	Database *DB* was created

METADATA_SERVICE_DB_DELETE

Severity	Description
informational	Database *DB* was deleted

IPINTERFACE_CREATE

Severity	Description
informational	A new iscsi IP interface was defined with name '*interface name*' on module *module* with port '*port list*' and IP address *IP address*

IPINTERFACE_DELETE

Severity	Description
informational	ISCSI IP interface with name '*interface name*' was deleted

IPINTERFACE_RENAME

Severity	Description
informational	ISCSI IP interface with name '*old name*' and was renamed '*interface name*'

IPINTERFACE_UPDATE

Severity	Description
informational	ISCSI IP interface with name '*interface name*' was updated. Its IP address is *IP address*

IPINTERFACE_UPDATE_MANAGEMENT

Severity	Description
informational	Management IP interfaces were updated. Management IPs are *IP addresses*

IPINTERFACE_UPDATE_MANAGEMENT_IPV6

Severity	Description
informational	Management IP interfaces were updated. Management IPv6 addresses are *IPv6 addresses*

IPINTERFACE_UPDATE_VPN

Severity	Description
informational	VPN IP interfaces were updated. VPN IPs are *IP addresses*

IPINTERFACE_UPDATE_VPN_IPV6

Severity	Description
informational	VPN IPv6 interfaces were updated. VPN IPv6 addresses are *IP addresses*

AUXILIARY_INTERNAL_PORTS_ENABLED

Severity	Description
informational	*Port Count* auxiliary internal Ethernet ports were enabled

AUXILIARY_INTERNAL_PORTS_DISABLED

Severity	Description
informational	*Port Count* auxiliary internal Ethernet ports were disabled

IPSEC_ENABLED

Severity	Description
informational	IPSec was enabled

IPSEC_DISABLED

Severity	Description
informational	IPSec was disabled

IPSEC_CONNECTION_ADDED

Severity	Description
informational	A new IPSec connection named '*name*' was added

IPSEC_CONNECTION_UPDATED

Severity	Description
informational	The IPSec connection named '*name*' was updated

IPSEC_CONNECTION_REMOVED

Severity	Description
informational	The IPSec connection named '*name*' was removed

PRIVATE_KEY_ADDED

Severity	Description
informational	A new private key named '*name*' with fingerprint '*fingerprint*' and size *key_size* bits was added.

CERTIFICATE_REMOVED

Severity	Description
informational	The certificate named '*name*' was removed.

PKCS12_CERTIFICATE_ADDED

Severity	Description
informational	A new PKCS#12 named '*name*' with fingerprint '*fingerprint*' was added.

PKI_RENAME

Severity	Description
informational	PKI with the name '*old name*' was renamed to '*new name*'

PKI_UPDATED

Severity	Description
informational	PKI with the name '*name*' and fingerprint '*fingerprint*' was updated

EMAIL_HAS_FAILED

Severity	Description	Troubleshooting
variable	Sending event *Event Code* (*Event Index*) to *Destination List* via *SMTP Gateway* failed. Module: *Module ID*; Error message: '*Error Message*'; timeout expired: *Timeout Expired?*.	Contact IBM Support

BULK_EMAIL_HAS_FAILED

Severity	Description	Troubleshooting
variable	Sending bulk email with *Events Number* events to *Destination List* via *SMTP Gateway* failed. Module: *Module ID*; Error message: '*Error Message*'; timeout expired: *Timeout Expired?*.	Contact IBM Support

SMS_HAS_FAILED

Severity	Description	Troubleshooting
variable	Sending event *Event Code* (*Event Index*) to *Destination List* via *SMS Gateway* and *SMTP Gateway* failed. Module: *Module ID*; Error message: '*Error Message*'; timeout expired: *Timeout Expired?*.	Contact IBM Support

HTTPS_HAS_FAILED

Severity	Description	Troubleshooting
variable	Sending event *Event Code* (*Event Index*) to *Destination List* via *HTTPS address* failed. Module: *Module ID*; Error message: '*Error Message*' (*HTTP error code*); timeout expired: *Timeout Expired?*.	Contact IBM Support

EMAIL_NOT_SENT

Severity	Description	Troubleshooting
variable	Sending event *Event Code* (*Event Index*) to *Destination List* via *SMTP Gateway* was waived because of failed SMTP gateway. It will be not be used until *Retry Time*.	Contact IBM Support

SMS_NOT_SENT

Severity	Description	Troubleshooting
variable	Sending event *Event Code* (*Event Index*) to *Destination List* via *SMS Gateway* and *SMTP Gateway* was waived because of failed SMTP gateway. It will be not be used until *Retry Time*.	Contact IBM Support

HEARTBEAT_EMAIL_HAS_FAILED

Severity	Description	Troubleshooting
minor	Sending heartbeat to *Destination Name* via *SMTP Gateway* failed. Module: *Module ID*; Error message: '*Error Message*'; timeout expired: *Timeout Expired?*.	Contact IBM Support

HEARTBEAT_SMS_HAS_FAILED

Severity	Description	Troubleshooting
minor	Sending heartbeat to *Destination Name* via *SMS Gateway* and *SMTP Gateway* failed. Module: *Module ID*; Error message: '*Error Message*'; timeout expired: *Timeout Expired?*.	Contact IBM Support

TEST_EMAIL_HAS_FAILED

Severity	Description	Troubleshooting
minor	Sending test to *Destination Name* via *SMTP Gateway* failed. Module: *Module ID*; Error message: '*Error Message*'; timeout expired: *Timeout Expired?*.	Contact IBM Support

TEST_SMS_HAS_FAILED

Severity	Description	Troubleshooting
minor	Sending test to *Destination Name* via *SMS Gateway* and *SMTP Gateway* failed. Module: *Module ID*; Error message: '*Error Message*'; timeout expired: *Timeout Expired?*.	Contact IBM Support

CUSTOM_EVENT

Severity	Description
variable	*Description*

UPGRADE_SOFTWARE_DOWNLOAD_FINISHED

Severity	Description
informational	Finished downloading software needed for upgrade to version *version*. Upgrade consequence is *consequence*

UPGRADE_FILE_LIST_RETRIEVAL_FAILED

Severity	Description	Troubleshooting
critical	Could not receive new version's file list from repository. Error code is *error*.	Contact IBM Support

UPGRADE_STARTS

Severity	Description
informational	System starting an upgrade.

PRE_UPGRADE

Severity	Description
informational	System preparing an upgrade procedure type *type* .

UPGRADE_IS_OVER

Severity	Description
informational	System went up after an upgrade.

IOS_RESTORED_AFTER_HOT_UPGRADE

Severity	Description
informational	System is able to perform I/Os after a hot upgrade.

UPGRADE_NO_NEW_FILES_FOR_UPGRADE

Severity	Description	Troubleshooting
warning	Repository version does not contain any new files. current version *current_version* new version is *new_version*	Contact IBM Support

UPGRADE_DOWNLOAD_REPOSITORY_COPY

Severity	Description	Troubleshooting
critical	Mirroring needed files from repository failed. Mirroring module is *mirroring_module* error is *error*	Contact IBM Support

UPGRADE_LOCAL_VERSION_DOWNLOAD_FAILED

Severity	Description	Troubleshooting
critical	Failure to distribute new sofware internally. Error code is *error* .	Contact IBM Support

UPGRADE_WAS_CANCELLED

Severity	Description	Troubleshooting
informational	Upgrade was cancelled with reason *reason* .	Contact IBM Support

HOT_UPGRADE_ABORTED

Severity	Description	Troubleshooting
critical	Hot upgrade aborted with reason *reason*.	Contact IBM Support

HOT_UPGRADE_HAS_FAILED

Severity	Description	Troubleshooting
critical	Hot upgrade failed while *errorneous_state*.	Contact IBM Support

PRE_UPGRADE_SCRIPT_INVOCATION_FAILED

Severity	Description	Troubleshooting
critical	Invocation of pre-upgrade script failed with error *error*.	Contact IBM Support

POST_UPGRADE_SCRIPT_INVOCATION_FAILED

Severity	Description	Troubleshooting
critical	Invocation of post-upgrade script failed with error *error*.	Contact IBM Support

UPGRADE_IS_NOT_ALLOWED

Severity	Description	Troubleshooting
critical	One or more of the pre-upgrade validations failed.	Fix the problems pointed out it previous events and revalidate.

PRE_UPGRADE_VALIDATION_FAILED

Severity	Description	Troubleshooting
critical	One of the pre-upgrade validations failed with status *error*.	Contact IBM Support

POST_UPGRADE_SCRIPT_STARTED

Severity	Description
informational	Post-upgrade script started.

POST_UPGRADE_SCRIPT_FINISHED

Severity	Description
informational	Post-upgrade script finished successfully.

PRE_UPGRADE_SCRIPT_DISAPPROVES

Severity	Description	Troubleshooting
critical	Upgrade cannot commence because some of the validations in the pre-upgrade script failed. Explanation: *explanation*.	Correct the system state according to the explanation and try again

POST_UPGRADE_SCRIPT_REPORTED_FAILURE

Severity	Description	Troubleshooting
critical	Post upgrade script reported failure. Script output: *explanation*.	Correct the system state according to the explanation and try again

POWER_PROBLEM_CAUSING_MAINTENANCE_MODE

Severity	Description
warning	Power state causing system to enter maintenance mode.

SYSTEM_ENTERED_CHARGING_STATE

Severity	Description
informational	System cannot start work until it is sufficiently charged.

POWER_PROBLEM_CAUSING_MODULE_PHASEOUT

Severity	Description
critical	Module *module* was phased out due to a power problem.

POWER_REPORT_PROBLEM_CAUSING_MODULE_PHASEOUT

Severity	Description
critical	Module *module* did not report power status on time and hence it was phased out.

POWER_PROBLEM_CAUSING_SYSTEM_SHUTDOWN

Severity	Description
critical	Power state causing system to shutdown due to: *Power Emergency Shutdown Reason*.

DELAYING_BACKUP_POWER_FAILURE_HANDLING

Severity	Description
informational	Delaying backup power failure handling at module *Module*

NO_DELAYED_BACKUP_POWER_FAILURE

Severity	Description
informational	No delayed backup power failure

MODULE_PHASEOUT_FAILED

Severity	Description
informational	Phase out of module *Module* failed

METADATA_SET

Severity	Description
warning	*Object type* with name '*Object name*' has new metadata value.

METADATA_DELETE

Severity	Description
warning	Metadata object deleted for *Object type* with name '*Object name*'.

SUBORDINATE_METADATA_SET

Severity	Description
warning	Remote *Object type* with name '*Object name*' was assigned a new metadata value by local system.

SUBORDINATE_METADATA_DELETE

Severity	Description
warning	Remote metadata object was deleted by local system for *Object type* with name '*Object name*'.

PATCH_SCRIPT_ADDED

Severity	Description	Troubleshooting
informational	Added patch *Patch Name*.	Was patch supposed to have been added.

PATCH_SCRIPT_UPDATED

Severity	Description
informational	Updated patch *Patch Name*.

PATCH_SCRIPT_DELETED

Severity	Description
informational	Deleted patch *Patch Name*.

MODULE_FAILED_TO_FETCH_PATCH_SCRIPT

Severity	Description
warning	Module *Module* failed to fetch patch script *Patch Name*.

PATCH_SCRIPT_FAILED_TO_EXECUTE

Severity	Description
informational	Patch script *Patch Name* execution failed on module *Module*

PATCH_SCRIPT_EXECUTION_STARTED

Severity	Description
informational	Patch script *Patch Name* execution on module *Module* started with pid *Process ID*

PATCH_SCRIPT_EXECUTION_ENDED

Severity	Description
informational	Patch script *Patch Name* execution on module *Module* with pid *Process ID* ended with return code *Return Code*

DOMAIN_CREATED

Severity	Description
informational	Domain *domain_name* has been created.

DOMAIN_UPDATED

Severity	Description
informational	Domain *domain_name* has been updated.

DOMAIN_RENAMED

Severity	Description
informational	Domain *old_name* has been renamed to *domain_name*.

DOMAIN_DELETED

Severity	Description
informational	Domain *domain_name* has been deleted.

POOL_ADDED_TO_DOMAIN

Severity	Description
informational	Pool *pool_name* has been added to domain *domain_name*.

POOL_REMOVED_FROM_DOMAIN

Severity	Description
informational	Pool *pool_name* has been removed from domain *domain_name*.

POOL_MOVED_BETWEEN_DOMAINS

Severity	Description
informational	Pool *pool_name* has been moved from domain *domain_name* to domain *domain_name*.

DOMAINS_AUTO_SHIFT_RESOURCES

Severity	Description
informational	Resources from domain *domain_name* to domain *domain_name* have been auto shifted.

OBJECT_ATTACHED_TO_DOMAIN

Severity	Description
informational	Object *object_name* of type *object_type* has been added to domain *domain_name*.

OBJECT_REMOVED_FROM_DOMAIN

Severity	Description
informational	Object *object_name* of type *object_type* has been removed from domain *domain_name*.

DOMAIN_MANAGED_ATTRIBUTE_SET

Severity	Description
informational	Domain *domain_name* managed attribute was set to *managed_attribute*.

REMOTE_SUPPORT_CONNECTED

Severity	Description
informational	System connected to remote support center *Destination*.

UNABLE_TO_CONNECT_TO_REMOTE_SUPPORT

Severity	Description
minor	System is unable to connect to any remote support center.

REMOTE_SUPPORT_CONNECTION_LOST

Severity	Description
variable	Connection to remote support center *Destination* failed while the connection was in state *Disconnected Session State*.

REMOTE_SUPPORT_TIMEOUT

Severity	Description
variable	Connection to remote support center *Destination* timed out while the connection was in state *Disconnected Session State*.

REMOTE_SUPPORT_IMMINENT_TIMEOUT

Severity	Description
minor	System is about to disconnect busy connection to remote support center *Destination*.

REMOTE_SUPPORT_DEFINED

Severity	Description
informational	Defined remote support center *Name* with IP address *Address* and port *Port*.

REMOTE_SUPPORT_DELETED

Severity	Description
informational	Deleted remote support center *Name*.

REMOTE_SUPPORT_DISCONNECTED

Severity	Description
variable	System disconnected from remote support center *Destination* while the connection was in state *Disconnected Session State*.

REMOTE_SUPPORT_CLIENT_MOVED

Severity	Description
informational	The remote support client moved from *Old Module* to *New Module*.

REMOTE_SUPPORT_CLIENT_NO_AVAILABLE_MODULES

Severity	Description
minor	No live modules with *Port Type* ports are available to run the remote support client.

TIMEZONE_SET

Severity	Description
informational	Timezone of the system was set to *Timezone*.

TIME_SET

Severity	Description	Troubleshooting
informational	On *Previous Time and Date* date and time of the system were set to *Time and Date*.	If date and/or time setting was intended, there is no problem.

TRANSACTION_ROLLED_BACK

Severity	Description
warning	Configuration transaction was rolled back due to module failure. Preceding events may reflect changes that were not committed.

ELICENSE_ACCEPTED

Severity	Description
informational	Electronic license was accepted by '*Approver Name*'.

ELICENSE_VIOLATION

Severity	Description	Troubleshooting
warning	Latest version of the electronic license was not approved.	Please approve the electronic license.

AUDIT_ENABLED

Severity	Description
informational	CLI command auditing activated.

AUDIT_DISABLED

Severity	Description
warning	CLI command auditing deactivated.

IB_PORT_ENABLE

Severity	Description
informational	Switch port *switch_port* has been enabled.

IB_PORT_DISABLE

Severity	Description
informational	Switch port *switch_port* has been disabled.

IB_PORT_MISWIRE

Severity	Description	Troubleshooting
warning	Switch port miswired: *switch_port* shall connect to *expected_component* but connects to *component* with GUID *guid*.	Contact IBM Support

IB_PORT_MISWIRE_CORRECTED

Severity	Description
informational	Miswire on switch port *switch_port* has been corrected.

IB_PORT_SHUTDOWN

Severity	Description	Troubleshooting
warning	Infiniband port *component* has been shutdown with reason *shutdown_reason*.	Contact IBM Support

IB_PORT_TEST_FAILED

Severity	Description	Troubleshooting
warning	Infiniband port *component* has failed component test with reason *reason*.	Contact IBM Support

IB_PORT_TEST_SUCCESS

Severity	Description	Troubleshooting
informational	Infiniband port *component* completes component test.	Contact IBM Support

IB_PORT_PHASEIN_FAILED

Severity	Description	Troubleshooting
warning	Infiniband port *component* has failed to phase-in with reason *reason*.	Contact IBM Support

IB_PORT_PHASEIN_SUCCESS

Severity	Description	Troubleshooting
informational	Infiniband port *component* completes phase-in.	Contact IBM Support

IB_LINK_DOWN

Severity	Description
warning	Link on *switch_port* (that connects to *component*) is down.

IB_LINK_UP

Severity	Description
informational	Link on *switch_port* (that connects to *component*) is up.

IB_SWITCH_MISSING

Severity	Description	Troubleshooting
warning	Switch *ib_switch* is missing.	Contact IBM Support

IB_SWITCH_LOST

Severity	Description	Troubleshooting
warning	Switch *ib_switch* that was missing is considered lost.	Contact IBM Support

IB_MISSING_SWITCH_FOUND

Severity	Description
informational	A previously missing switch *ib_switch* is now found.

IB_CONFIGURE_COMMAND_ERROR

Severity	Description
warning	Infiniband configuration command *command* has failed on *component*.

IB_PERF_COUNTER_RESET

Severity	Description
informational	Performance counter *counter* will be reset on *component*.

POD_IB_PORT_MISWIRE

Severity	Description	Troubleshooting
warning	POD module port miswired: *module_port* shall connect to *expected_component* but connects to *component* with GUID *guid*.	Please contact support.

POD_IB_PORT_MISWIRE_CORRECTED

Severity	Description
informational	POD module port *module_port* connected to *component* miswire corrected.

IB_SWITCH_REBOOT_DETECTED

Severity	Description	Troubleshooting
warning	IB switch '*switch_id*' has rebooted.	Contact IBM Support

PERF_CLASS_MAX_IO_RATE_UPDATED

Severity	Description
informational	Performance Class *name* max IO rate was changed to *IO rate* IOPS

PERF_CLASS_MAX_BW_RATE_UPDATED

Severity	Description
informational	Performance Class *name* max BW rate was changed to *BW rate* MBps

PERF_CLASS_CREATE

Severity	Description
informational	Performance Class with name '*name*' was created

PERF_CLASS_DELETE

Severity	Description
informational	Performance Class with name '*name*' was deleted

PERF_CLASS_ADD_HOST

Severity	Description
informational	Host with name '*host_name*' was added to Performance Class with name '*name*'

PERF_CLASS_REMOVE_HOST

Severity	Description
informational	Host with name '*host_name*' was removed from Performance Class with name '*name*'

PERF_CLASS_ADD_POOL

Severity	Description
informational	Pool with name '*pool.name*' was added to Performance Class with name '*pool.perf_class*'

PERF_CLASS_REMOVE_POOL

Severity	Description
informational	Pool with name '*pool.name*' was removed from Performance Class with name '*name*'

PERF_CLASS_ADD_VOLUME

Severity	Description
informational	Volume with name '*volume.name*' was added to Performance Class with name '*volume.perf_class*'

PERF_CLASS_REMOVE_VOLUME

Severity	Description
informational	Volume with name '*volume.name*' was removed from Performance Class with name '*name*'

PERF_CLASS_ADD_DOMAIN

Severity	Description
informational	Domain *domain_name* was added to Performance Class *name*

PERF_CLASS_REMOVE_DOMAIN

Severity	Description
informational	Domain *domain_name* was removed from Performance Class *name*

VOLUME_MODIFIED_DURING_IO_PAUSE

Severity	Description	Troubleshooting
warning	Volume '*vol_name*' of CG '*cg_name*' was modified during Pause IO with token '*token*'	Retry after completing CG changes.

CONS_GROUP_MODIFIED_DURING_IO_PAUSE

Severity	Description	Troubleshooting
warning	CG '*cg_name*' was modified during Pause IO with token '*token*'	Retry after completing CG changes.

IO_PAUSED_FOR_CONS_GROUP

Severity	Description
Informational	Pause IO on CG with name '*cg_name*' was started with *timeout*ms timeout . Token is '*token*'.

IO_RESUMED_FOR_CONS_GROUP_EXPLICITLY

Severity	Description
Informational	Pause IO on CG with name '*cg_name*' and token '*token*' was resumed by user request.

IO_RESUMED_FOR_CONS_GROUP_AUTOMATICALLY

Severity	Description
Informational	Pause IO on CG with name '*cg_name*' and token '*token*' was resumed after snapgroup creation.

IO_RESUMED_FOR_CONS_GROUP_UPON_SYSTEM_ERROR

Severity	Description
warning	Pause IO on CG with name '*cg_name*' and token '*token*' was resumed after system error.

IO_RESUMED_FOR_CONS_GROUP_ UPON_TIMEOUT_EXPIRATION

Severity	Description	Troubleshooting
warning	Pause IO on CG with name '*cg_name*' and token '*token*' was canceled after timeout.	Use longer timeout value or require less time for performing action.

ALU_CREATE

Severity	Description
informational	ALU was defined with name '*ALU name*' associated with host '*ALU host name*' lun '*ALU lun*'.

ALU_DELETE

Severity	Description
informational	ALU with name '*ALU name*' associated with host '*ALU host name*' lun '*ALU lun*' was deleted.

ALU_UNBOUND_ALL

Severity	Description
informational	All SLUs of ALU with name '*ALU name*' were unbound.

TXN_NODE_FLASH_CONNECTION_LOST

Severity	Description
warning	TXN node on '*module*' is disconnected from flash system '*flash system*'.

TXN_NODE_FLASH_CONNECTED

Severity	Description
informational	TXN node on '*module*' is fully connected to flash system '*flash system*'.

POD_IB_LINK_DETECTION_LINK_ PERSISTENTLY_DISCONNECTED

Severity	Description
critical	IB link from *Source* to *Target* has reported as persistently disconnected

POD_IB_LINK_DETECTION_LINK_PERSISTENTLY_CONNECTED

Severity	Description
informational	IB link from *Source* to *Target* has detected as persistently connected

Chapter 26. 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.

© Copyright IBM Corp. 2017

Notices

These legal notices pertain to the information in this IBM Storage product documentation.

This information was developed for products and services offered in the US. This material may be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
USA

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

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

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some 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.

© Copyright IBM Corp. 2017

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 Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
USA

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.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

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

Trademarks

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

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



Printed in USA

SC27-8559-06

