IBM FlashSystem A9000
12.3.2.e

# *Command-Line Interface (CLI) Reference Guide*

**Note**

Before using this document and the product it supports, read the information in “Notices” on page 741.

**First Edition (January 2023)**

This edition applies to IBM FlashSystem® A9000 12.3.2.e. Newer document editions may be issued for the same product version in order to add missing information, update information, or amend typographical errors. The edition is reset to 'First Edition' for every new product version.

**© Copyright International Business Machines Corporation 2016, 2022.**
US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

# Contents


**About this guide.....xiii**
Intended audience..... xiii
Conventions used in this guide.....xiii
Related information and publications.....xiii
IBM Publications Center.....xiii
Feedback.....xiv
Getting information, help, and service.....xiv

**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..... 3
Interactive and basic modes.....4
Understanding the command syntax..... 5
Using identification parameters.....7
Display options..... 10
Running commands in batch mode..... 12
Displaying CLI help and the XCLI utility version.....13

**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..... 19
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, clusters, or domains..... 33
Setting the special type of hosts or clusters..... 34
Listing hosts/clusters to which a volume is mapped..... 36
Unmapping a volume from a host or cluster..... 36
Set the default idle time before unmapping a volume..... 38
Retrieving the default idle time before unmapping a volume..... 39
Creating a performance class.....39
Deleting a performance class.....40
Renaming a performance class..... 41
Listing details of performance classes.....42
Adding a host to a performance class..... 43
Removing a host from its performance class.....44
Adding a pool to a performance class..... 45
Removing a pool from its performance class.....46

Adding a volume to a performance class....47
Removing a volume from its performance class....48
Adding a domain to a performance class....48
Removing a domain from its performance class....49
Setting the rate for a performance class....50
Listing host profiles....51
Updating the host profile....52
Removing the profile of the specified host....53
Enabling the host profiler....53
Disabling the host profiler....54

**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....66
Listing a volume's extended attributes....68
Locking a volume....69
Renaming a volume....71
Resizing a volume....73
Unlocking a volume....76

**Chapter 4. Volume snapshot management commands ....79**
Changing a snapshot deletion priority....79
Creating a snapshot....80
Deleting a snapshot....83
Duplicating a snapshot....85
Formatting a snapshot....86
Listing snapshot information....87
Restoring a volume from a snapshot....89

**Chapter 5. Consistency group management commands ....93**
Adding a volume to a consistency group....93
Creating consistency groups....97
Deleting a consistency group....99
Listing consistency groups....100
Removing a volume from a consistency group....101
Renaming a consistency group....103
Creating a cross-system consistency group....104
Associating an existing consistency group with a cross-system consistency group definition....105
Removing an existing consistency group from a cross-system consistency group definition....106
Adding a remote system name to a cross-system consistency group definition....107
Removing a remote system name from a cross-system consistency group definition....108
Listing cross-system consistency group definitions....109
Retrieving remote systems in a specified cross-system consistency group....110
Deleting a cross-system consistency group....111
Listing cross-system consistency group definitions....112

**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....120
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....128

**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....140
Resizing a storage pool....141
Moving a volume between storage pools....143

**Chapter 8. System management commands....147**
Displaying the values of configuration parameters....147
Setting configuration parameters....149
Testing the DNS....150
Displaying help....151
Printing the current maintenance urgency....152
Shutting down the system....153
Listing the operational state....154
Displaying system usage and data reduction statistics....155
Displaying information about physical and effective capacity....156
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....163
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....171
Enabling command auditing....172
Disabling command auditing....172
Displaying the command audit state....173
Configuring audit servers....174
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....199

**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....208
Changing the RPO for local or remote system....211
Changing the designation of mirroring peers....213
Changing a mirroring schedule for remote slave peers....215
Changing the role of a mirrored volume....217
Changing a mirroring schedule for local peers....220
Creating a mirroring definition....222
Deactivating mirroring....229
Deleting a remote mirroring definition....231
Viewing the mirroring status....234
Obtaining statistics on past sync jobs....237
Switching roles between master and slave....239
Retrieving RPO thresholds....241
Setting an RPO threshold....241
Changing the interval of a schedule....242
Creating a schedule object....244
Triggering a schedule....246
Deleting a schedule object....247
Listing a schedule object....248
Renaming a schedule....249
Viewing sync job status....249

**Chapter 11. HyperSwap commands....251**
Creating a HyperSwap relation....251
Viewing the status of HyperSwap volumes and consistency groups....258
Activating a HyperSwap relation....260
Deactivating a HyperSwap relation....263
Deleting a HyperSwap relation....265
Switching roles between Master and Slave volumes....268
Changing a peer role in a HyperSwap volume....270
Restoring the availability of a Master volume....273
Creating a HyperSwap volume snapshot (ad hoc sync job)....274
Changing the designation of HyperSwap relation peers....278

Enabling high availability of a HyperSwap relation....280
Disabling high availability of a HyperSwap relation....281
Converting a HyperSwap relation into a sync mirror....282
Converting a sync mirror into a HyperSwap relation....284
Creating a new Quorum Witness....287
Listing Quorum Witnesses....289
Updating a Quorum Witness definition....291
Renaming a Quorum Witness....292
Deleting a Quorum Witness....293
Activating a Quorum Witness....294
Deactivating a Quorum Witness....295
Listing the Quorum Witness connection status....296
Getting Quorum Witness information....298

**Chapter 12. Data migration commands ....301**
Activating data migration....301
Deactivating data migration....302
Defining data migration configuration....303
Deleting a data migration process....305
Listing data migration statuses....307
Testing the data migration definition....308

**Chapter 13. IBM Hyper-Scale Mobility commands ....311**
Creating an IBM Hyper-Scale Mobility relation....311
Activating a volume migration....315
Deactivating IBM Hyper-Scale Mobility migration....317
Aborting a defined or activated IBM Hyper-Scale Mobility process....318
Moving the IBM Hyper-Scale Mobility source volume to a Proxy state....319
Deleting an IBM Hyper-Scale Mobility relation....321
Listing the IBM Hyper-Scale Mobility status....322

**Chapter 14. Event handling commands ....325**
Generating a custom event....325
Generating a CSS product event....325
Defining a new event notification destination....326
Deleting a destination....330
Listing event notification destinations....331
Renaming a destination....332
Testing a destination....333
Updating an event notification destination....334
Adding a destination to a destination group....338
Creating a destination group....339
Updating an event notification destination group....340
Deleting a destination group....341
Listing destination groups....342
Removing a destination from a destination group....342
Renaming a destination group....343
Clearing alerting events....344
Listing events....345
Listing uncleared alerting events....348
Setting the threshold for event notification....348
Listing thresholds....349
Activating a rule....351
Creating event notification rules....351
Deactivating a rule....354
Deleting event notification rules....355
Listing event notification rules....356

Renaming event notification rules..... 357
Updating an event notification rule..... 358
Defining an SMS gateway..... 360
Deleting an SMS gateway..... 362
Listing SMS gateways..... 363
Prioritizing SMS gateways..... 364
Renaming an SMS gateway..... 365
Updating an SMS gateway..... 366
Defining a new SMTP gateway..... 367
Deleting an SMTP gateway..... 368
Listing SMTP gateways..... 369
Prioritizing SMTP gateways..... 370
Renaming an SMTP gateway..... 371
Updating an SMTP gateway..... 372
Generating an XMPNS admin control event..... 374

**Chapter 15. IP configuration commands ..... 375**
Updating the configuration of the Ethernet port..... 375
Showing the status and configuration of Ethernet ports (deprecated)..... 376
Showing the status and configuration of Ethernet ports..... 378
Displaying Ethernet (physical) port statistics in the system..... 380
Creating a new IP interface..... 382
Deleting IP interfaces..... 384
Listing IP interface configuration..... 385
Listing IP interface addresses..... 386
Displaying IP interface statistics in the system..... 387
Renaming an IP interface..... 388
Printing the ARP database of an IP interface..... 389
Testing the traceroute to a remote IP..... 390
Testing the traceroute to a remote IPv6 address..... 391
Updating an IP interface..... 391
Updating VLAN priority code points..... 394
Defining a new IPSec connection..... 395
Updating an existing IPSec connection..... 397
Removing an existing IPSec connection..... 398
Listing IPSec connections..... 399
Listing IPSec tunnels..... 400
Connecting to a support center..... 400
Defining a support center..... 402
Deleting a support center..... 403
Disconnecting from a support center..... 403
Listing support centers..... 404
Listing statuses of all support centers..... 405
Configuring the support center connection to enable automatic connect on restart..... 406
Listing the configuration of support center automatically connect..... 407
Creating a new IP access group..... 408
Removing an address from an IP access group..... 408
Adding a new address to an IP access group..... 409
Deleting an existing IP access group..... 410
Renaming an existing IP access group..... 411
Listing IP access groups..... 412
Listing IP access group addresses..... 413

**Chapter 16. PKI configuration commands ..... 415**
Listing PKI items..... 415
Generating a certificate signing request..... 416
Generating a private key and CSR..... 416

Deleting a PKI content....418
Changing a PKI symbolic name....418
Importing a signed certificate....419
Importing a PKCS#12 certificate....421
Displaying the details of a signed certificate....422
Updating a PKI certificate or services....423

**Chapter 17. InfiniBand commands ....425**
Listing the configured InfiniBand ports....425
Listing data counters for the enabled InfiniBand HCA ports....426
Listing error counters for the enabled InfiniBand HCA ports....427
Listing the statuses of the enabled InfiniBand HCA ports....428
Listing the configured InfiniBand switches....429
Listing the configured InfiniBand switch management addresses....431
Listing the configured InfiniBand switch firmware versions....432
Listing the configured InfiniBand switch power values....433
Listing the configured InfiniBand switch voltage values....434
Listing the configured InfiniBand switch temperature values....436
Listing the configured InfiniBand switch fan parts....437
Listing the configured InfiniBand switch PSUs....438
Listing the configured InfiniBand switch BBUs....439
Listing the configured InfiniBand switch fans....441

**Chapter 18. Access control commands ....443**
Adding an access control definition....443
Deleting an access control definition....444
Listing access control definitions....445
Adding an LDAP server definition....446
Testing an LDAP configuration....447
Listing LDAP configuration parameters....449
Configuring LDAP in the system....450
Listing LDAP servers defined in the system....454
Listing LDAP server users....455
Listing LDAP-based authentication mode....456
Enabling or disabling LDAP-based authentication mode....457
Updating an LDAP server definition....458
Removing an LDAP server definition....460
Launching the ldapsearch utility....461
Defining a new user....462
Deleting a user....465
Adding users to user groups....466
Creating user groups....467
Deleting a user group....468
Listing user groups....469
Removing a user from a user group....470
Renaming user groups....471
Updating a user group....472
Listing users....473
Renaming users....475
Updating a user definition....476
Creating a new domain....478
Updating a domain definition....480
Renaming a domain....482
Deleting a domain....483
Listing domains....484
Listing users per domain....486
Listing objects in domains....487

Listing the global domain.... 488
Attaching an object to a domain....489
Disassociating object from a domain.... 490
Associating users to a domain....493
Removing a user from a domain....494
Adding a pool to a domain....495
Removing a pool from a domain....497
Moving a pool from one domain to another.... 498
Setting the domain attribute.... 500
Setting domain related policies....501
Displaying domain-related policies....501
Specifying a user associated with IBM Hyper-Scale Manager.... 503
Retrieving the user associated with the IBM Hyper-Scale Manager....503
Setting the application administrator's scope of commands....504
Getting the application administrator's scope of commands.... 505

**Chapter 19. Fibre Channel and iSCSI configuration and status commands ..... 507**
Discovering FC hosts.... 507
Changing FC port configuration....508
Listing FC ports.... 509
Listing FC Port Tests.... 511
Starting FC Port Test....513
Aborting FC Port Test....514
Listing connectivity to hosts....515
Displaying the InfiniBand connectivity status.... 516

**Chapter 20. Flash enclosure maintenance commands ..... 519**
Listing Flash enclosure status....519
Listing a Flash canister status.... 520
Listing a Flash card status.... 522
Retrieving the list of Flash fans.... 523
Retrieving the list of Flash IB adapters....524
Retrieving the Flash control connectivity list....525
Retrieving the list of Flash PSUs....527
Retrieving the list of Flash PIBs.... 528
Retrieving the list of Flash LED cards....529
Listing Flash BBU components....530
Retrieving the serial number of a flash enclosure 1S.... 531
Suspending the automatic Flash BBU calibration.... 532

**Chapter 21. Hardware maintenance commands ..... 533**
Monitoring the rebuilding or redistribution processes.... 533
Displaying the system's average power consumption....534
Getting the values for calculating the system's average power consumption.... 534
Displaying the system's average temperature....535
Enabling XIV Support access.... 536
Disabling XIV Support access.... 537
Displaying the XIV Support window....538
Listing system components....539
Listing module configuration....542
Listing the internal temperature of modules.... 544
Listing boot media devices in the system.... 546
Listing vault devices in the system....548
Listing BBUs in the system.... 550
Listing PSUs in the system....553
Listing compression adapters in the system....554
Listing fans in a module.... 555

Listing NICs in the system.....557
Listing DIMMs in the modules.....559
Listing CPUs in the modules.....563
Listing InfiniBand HCA Adapters in the storage system.....565
Listing CNA adapters in the system.....566
Listing module LEDs in the system.....567
Listing data disk devices in the system.....569
Listing service statuses.....571
Listing system components that require service.....572
Listing trace snapshot on a module.....574
Creating a trace snapshot.....574

**Chapter 22. Statistics commands ..... 577**
Getting performance statistics.....577
Retrieving usage history.....582

**Chapter 23. Metadata commands ..... 585**
Setting metadata.....585
Deleting metadata.....586
Listing metadata.....587
Setting user metadata.....588
Listing user metadata.....589
Deleting user metadata.....590

**Chapter 24. Encryption enablement and support commands ..... 591**
Disabling encryption.....591
Enabling encryption.....593
Defining a key server.....595
Removing a key server.....597
Displaying key server status.....598
Checking key server status.....599
Obtaining a new master key.....600
Renaming a key server.....601
Changing key server properties.....602
Entering a recovery key.....604
Generating recovery keys.....606
Retrieving the security administrator's recovery key.....608
Rekeying the security administrators.....609
Displaying recovery key status.....611
Recovering key verification.....612
Recovering key share information.....613
Finishing the recovery process.....614
Obtaining a new master key.....615
Changing the key management scheme.....617
Viewing the key scheme.....618

**Chapter 25. Multi-site high availability and disaster recovery commands .....621**
Viewing the Multi-site relation status.....621
Creating a Multi-site relation.....624
Activating an asynchronous mirror within a Multi-site relation.....627
Deleting a Multi-site relation.....630
Changing roles in a Multi-site relation.....631
Adding a Standby mirror to a Multi-site relation.....634
Switching roles between Master and SMaster.....636

**Chapter 26. Security configuration commands ..... 639**
Listing configuration parameters for a communication protocol.....639

Setting configuration parameters for a communication protocol..... 640

**Chapter 27. Events..... 641**

**Chapter 28. Return Codes..... 739**

**Notices..... 741**

Trademarks..... 742

# 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 FlashSystem A9000 on the following information sources:

- IBM FlashSystem A9000 documentation (ibm.com®/docs/en/flashsystem-a9000) – where 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 documentation (ibm.com/docs/en/hyper-scale-manager)

## 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.

© Copyright IBM Corp. 2016, 2022

## Sending comments

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

### Procedure

To submit any comments about this publication or any other IBM storage product documentation:

- Send your comments by email to ibmdocs@us.ibm.com. Be sure to include the following information:
  - Exact publication title and version
  - Publication form number (for example, GA32-1234-00)
  - Page, table, or illustration numbers that you are commenting on
  - A detailed description of any information that should be changed

## 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 (ibm.com/mysupport)
- IBM Directory of Worldwide Contacts website (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" on page 1
- "Using the CLI" on page 3


## 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" on page 1
- "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.

#### About this task

**Note:** For the installation requirements and a list of available packages, see the *IBM XCLI Utility Release Notes* on IBM FlashSystem A9000 documentation (ibm.com/docs/en/flashsystem-a9000).

#### Procedure

Perform these steps to download and install the XCLI:

1. Download the IBM XCLI installation package from the IBM Fix Central website.
2. Perform one of the following procedures for your operating system.

   **Note:** See the IBM XCLI Utility release notes for your version to see which operating systems are supported. (See **Release notes** > **Release note PDF publications** on IBM Documentation).

   - **Windows:** Double-click the installation file, and follow the instructions on the screen.
   - **Linux®, AIX®, HP-UX, Solaris:** Run the installation file using the following command:

     ```
     ./<filename>
     ```

     Then set the program path.
3. Start the XCLI depending on the hosting operating system and operational mode.

### Starting the XCLI on a Windows system

You can start the XCLI on a Windows system in either interactive or basic mode.

© Copyright IBM Corp. 2016, 2022

### *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

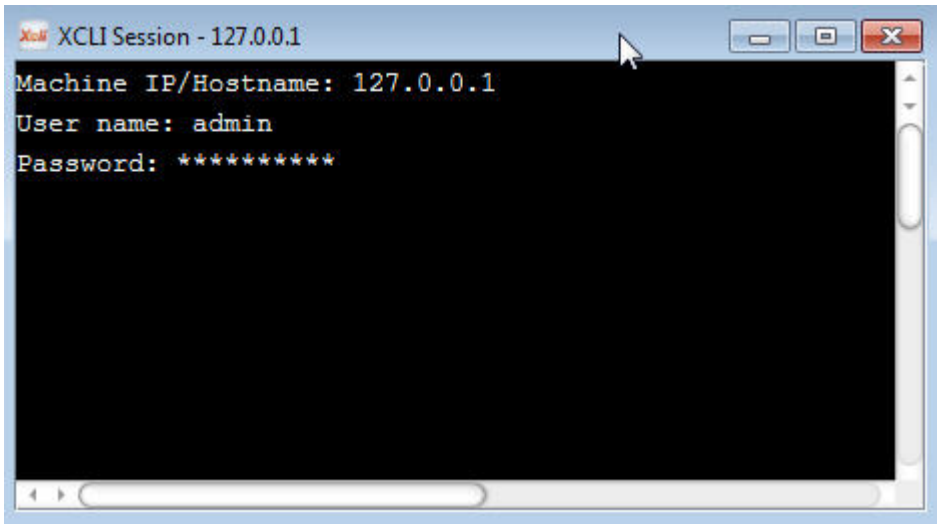


*Figure 1. XCLI terminal*

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" on page 4
- "Understanding the command syntax" on page 5
- "Using identification parameters" on page 7
- "Display options" on page 10
- "Running commands in batch mode" on page 12
- "Displaying CLI help and the XCLI utility version" on page 13

# 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 |

| Options | Values |
|---|---|
| `-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" on page 7
- "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" on page 10
- "Table display options" on page 11


### 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 overriden with the cluster's mapping and type.
- If **map=host**, the mapping of the cluster and its host type are overridden with the host's mapping and type. Use this value to add a host to an empty cluster, so that the cluster will obtain the host’s mapping.
- If **map=clusterWithHostExceptions**, the host keeps its mapping and the cluster mapping is appended on top of it.

The host or cluster receives a single SCSI unit attention message, even if the change affects multiple volumes.

**Example:**

```
cluster_add_host cluster=Cluster1 host=Host1 map=cluster
```

**Output:**

```
Command completed successfully.
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Allowed |
| Application administrator | Disallowed |
| Security administrator | Disallowed |

© Copyright IBM Corp. 2016, 2022

| User Category | Permission |
|---|---|
| Read-only users | Disallowed |
| Technicians | Disallowed |

### Return codes

- **HOST_BAD_NAME**

  The host name does not exist.
- **CLUSTER_BAD_NAME**

  The cluster name does not exist.
- **HOST_BELONGS_TO_ANOTHER_CLUSTER**

  This host already belongs to another cluster.
- **HOST_AND_CLUSTER_HAVE_CONFLICTING_MAPPINGS**

  Host mapping conflicts with cluster mapping.
- **HOST_AND_CLUSTER_HAVE_DIFFERENT_MAPPING_TYPE**

  The host mapping type is not the same as the cluster mapping type.
- **HOST_NOT_IN_CLUSTERS_DOMAINS**

  The host is not part of all of the domains the cluster is attached to.

## Creating a cluster

Use the **cluster_create** command to create a new cluster.

```
cluster_create cluster=ClusterName [ domain=DomainList ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **cluster** | Object name | Name of the cluster to be created. | Y | N/A |
| **domain** | N/A | The cluster will be attached to the specified domains. To define more than one domain, separate them with a comma. To attach the cluster to all existing domains, use "*". | N | none |

The newly created cluster does not contain hosts, and has the default type, but no mapping.

**Example:**

```
cluster_create cluster=Cluster1
```

**Output:**

```
Command completed successfully.
```

### Access control

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

### Return codes

- **CLUSTER_NAME_EXISTS**

  The cluster name already exists.
- **MAX_CLUSTERS_REACHED**

  The maximum number of defined clusters is already reached.
- **DOMAIN_DOESNT_EXIST**

  The domain does not exist.

## Deleting clusters

Use the **cluster_delete** command to delete a cluster.

```
cluster_delete cluster=ClusterName
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **cluster** | Object name | Cluster to be deleted. | Y |

This command deletes a cluster. All hosts contained in the cluster remain active and are not deleted. The special type of each host is set to the cluster's special type. The mapping of each host is set to the cluster's mapping. No I/O interruption is caused by this command.

**Example:**

```
cluster_delete cluster=Cluster1
```

**Output:**

```
Command completed successfully.
```

### Access control

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

| User Category | Permission |
|---|---|
| Technicians | Disallowed |

### Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_DELETE_ASSIGNED_CLUSTER**

  Cluster *Cluster*' includes hosts. Are you sure you want to delete it?

### Return codes

- **CLUSTER_BAD_NAME**

  The cluster name does not exist.

## Listing clusters

Use the **cluster_list** command to retrieve information about a specific cluster, or about all of them.

```
cluster_list [ cluster=ClusterName ] [ domain=DomainName ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **cluster** | Object name | Name of cluster to be listed. | N | All clusters. |
| **domain** | Object name | The domain name. | N | All Domains |

The output provides each cluster's special type, and comma-separated lists of hosts, users, and user groups.

**Example:**

```
cluster_list
```

**Output:**

```
Name       Hosts   Type      Creator           User Group
---------- ------- --------- ----------------- ------------
Cluster1           default   xiv_maintenance
```

| Field ID | Field output | Default position |
|---|---|---|
| **name** | Name | 1 |
| **hosts** | Hosts | 2 |
| **type** | Type | 3 |
| **creator** | Creator | 4 |
| **user_group** | User Group | 5 |

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Allowed |

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

# Removing a host from a cluster

Use the **cluster_remove_host** command to remove a host from a cluster.

```
cluster_remove_host cluster=ClusterName host=HostName
```

## Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **cluster** | Object name | Cluster name. | Y |
| **host** | Object name | Host to be removed from cluster. | Y |

This command removes the specified host from a cluster. The host then no longer belongs to any cluster. The host's special type and mapping remain identical to the cluster's special type and mapping, and therefore, I/O is not interrupted. The association of the host with user or user groups remains the same as the cluster's association.

## Access control

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

## Return codes

- **HOST_BAD_NAME**

  The host name does not exist.
- **CLUSTER_BAD_NAME**

  The cluster name does not exist.
- **HOST_NOT_IN_CLUSTER**

  This host does not belong to the specified cluster.

# Renaming clusters

Use the **cluster_rename** command to rename a cluster.

```
cluster_rename cluster=ClusterName new_name=Name
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **cluster** | Object name | Cluster to be renamed. | Y |
| **new_name** | Object name | New name of cluster. | Y |

This command renames the specified cluster.

**Example:**

```
cluster_rename cluster=Cluster1 new_name=Cluster2
```

**Output:**

```
Command completed successfully.
```

### Access control

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

### Return codes

- **CLUSTER_BAD_NAME**

  The cluster name does not exist.
- **CLUSTER_NAME_EXISTS**

  The cluster name already exists.

## Adding a port to a host

Use the **host_add_port** command to add a port address to a host.

```
host_add_port host=HostName < fcaddress=wwpn | iscsi_name=iSCSIName [ num_of_visible_targets=num
| num_of_visible_targets_per_network=num ] >
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **host** | Object name | Host name. | Y | N/A |
| **fcaddress** | N/A | FC address of the added port. | N | N/A |
| **iscsi_name** | iSCSI initiator name | iSCSI initiator name of the newly added port. | N | N/A |

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **num_of_visible_targets** | Integer | Limit the maximum number of target IP addresses which will be reported on iSCSI discoveries invoked by this initiator port. Valid values: 0 (unlimited), 2-64. | N | 0 |
| **num_of_visible_targets_per_network** | Integer | Maximum number of target IP addresses to be reported on iSCSI discoveries invoked by this initiator port per network. Valid values: 0 (unlimited), 2-64. | N | 0 |

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

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

**Example:**

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

**Output:**

```
Command completed successfully.
```

## Access control

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

## Return codes

- **HOST_BAD_NAME**

  The host name does not exist.
- **HOST_PORT_EXISTS**

  A host with this port ID is already defined.
- **ISCSI_HOST_ILLEGAL_PORT_NAME**

  The port name for iSCSI Host is illegal.

  **Troubleshooting:** Port names for iSCSI Hosts must contain only printable characters.
- **MAX_PORTS_REACHED**

  The maximum number of ports defined in the system is already reached.

- **TARGET_PORT_BAD_ADDRESS**

  The remote port address is illegal or does not belong to the remote target.
- **PORT_EXISTS**

  The port is already defined.
- **OLVM_LINK_IS_NOT_UP**

  The IBM Hyper-Scale Mobility link is not up. The mapping list cannot be updated.
- **TARGET_NOT_CONNECTED**

  There is currently no connection to the target system.
- **REMOTE_MAX_VIRTUAL_HOSTS_REACHED**

  The maximum number of defined remote virtual hosts is already reached.
- **INVALID_NUM_OF_TARGETS**

  The specified value representing the number of visible targets is invalid. It can be either 0 (unlimited), or a number in the range from 2 through 64.
- **PARAMETER_CANNOT_BE_UPDATED_ON_SYSTEM_HAVING_IPINTERFACE_WITH_VLAN**

  The num_of_visible_targets parameter cannot be used in a system with an IP interface that is configured for VLAN.

# Defining a new host

Use the **host_define** command to define a new host to connect to the storage system.

```
host_define host=HostName [ cluster=ClusterName ]
[ iscsi_chap_name=iscsiChapName iscsi_chap_secret=iscsiPass ] [ domain=DomainList ]
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **host** | Object name | The name of the host to be created. | Y | N/A |
| **cluster** | Object name | The name of the cluster to contain the newly created host. | N | No cluster. |
| **iscsi_chap_name** | String | The host's CHAP name identifier. | N | [None] |
| **iscsi_chap_secret** | String | The password of the initiator used to authenticate to the system when CHAP is enabled. | N | [None] |
| **domain** | N/A | The domains the cluster will be attached to.<br>To include more than one domain, separate them with a comma.<br>To include all existing domains, use an asterisk ("*"). | N | none |

The name of the host must be unique in the system.

Use the **host_add_port** command to add port addresses to this host (see Adding a port to a host for details). Specifying the cluster is optional.

The parameters **iscsi_chap_name** and **iscsi_chap_secret** must be either both specified or both unspecified.

If **iscsi_chap_secret** does not conform to the required secret length (96-128 bits), the command will fail.

The command checks whether the **iscsi_chap_name** and **iscsi_chap secret** are unique. In case they are not, an error message is displayed, but the command completes.

The secret has to be between 96 bits and 128 bits. There are 3 ways to enter the secret:

- *Base64*: Requires to prefix the entry with 0b. Each subsequent character entered is treated as a 6-bit equivalent length
- *Hex*: Requires to prefix the entry with 0x. Each subsequent character entered is treated as a 4-bit equivalent length
- *String*: Requires no prefix (cannot be prefixed with 0b or 0x). Each character entered is treated as a 8 bit equivalent length

**Example:**

```
host_define host=server1
```

**Output:**

```
Command completed successfully.
```

## Access control

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

## Warnings

- **ISCSI_CHAP_NAME_AND_SECRET_NOT_UNIQUE**

  Both iSCSI CHAP name and secret are already used by another host. Are you sure you want to reuse those values?

## Return codes

- **HOST_NAME_EXISTS**

  The host name already exists.
- **MAX_HOSTS_REACHED**

  The maximum number of defined hosts is already reached.
- **CLUSTER_BAD_NAME**

  The cluster name does not exist.
- **DOMAIN_DOESNT_EXIST**

The domain does not exist.

- **HOST_NOT_IN_CLUSTERS_DOMAINS**

  The host is not part of all of the domains the cluster is attached to.

# Deleting a host

Use the **host_delete** command to delete a host.

```
host_delete host=HostName
```

## Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **host** | Object name | The host name. | Y |

After this command is executed, the deleted host can no longer connect to the system, and I/O requests from this host are not handled.

**Example:**

```
host_delete host=mailserver
```

**Output:**

```
Command completed successfully
```

## Access control

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

## Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_DELETE_HOST**

  Are you sure you want to delete host *Host*?

## Return codes

- **HOST_BAD_NAME**

  The host name does not exist.

# Listing hosts

Use the **host_list** command to list a specific host or all hosts.

```
host_list [ host=HostName ] [ perf_class=perfClassName ] [ domain=DomainName ]
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **host** | Object name | The host name. | N | All hosts. |
| **perf_class** | Object name | The name of a performance class. | N | no filter. |
| **domain** | Object name | The domain name. | N | All Domains |

This command lists all the hosts in the system.

A host name can be specified to list only a specific host or all the hosts.

The list contains the following comma separated information:

- Port addresses
- Containing cluster, if one exists
- Associated users and user groups

**Example:**

```
host_list host=mailserver
```

**Output:**

```
Name     Type      FC Ports   iSCSI Ports   User Group   Cluster
-------- --------- ---------- ------------- ------------ ---------
host_4   default              iscsi_4
host_5   default              iscsi_5
host_6   default              iscsi_6
host_7   default              iscsi_7
host_8   default              iscsi_8
host_9   default              iscsi_9
```

| Field ID | Field output | Default position |
|---|---|---|
| **name** | Name | 1 |
| **type** | Type | 2 |
| **fc_ports** | FC Ports | 3 |
| **iscsi_ports** | iSCSI Ports | 4 |
| **creator** | Creator | N/A |
| **user_group** | User Group | 5 |
| **cluster** | Cluster | 6 |
| **perf_class** | Performance Class | 7 |
| **iscsi_chap_name** | iSCSI CHAP Name | N/A |

### Access control

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

## Listing ports

Use the **host_list_ports** command to list all the ports of a host.

```
host_list_ports host=HostName [ domain=DomainName ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **host** | Object name | The host name. | Y | N/A |
| **domain** | Object name | The domain name. | N | All Domains |

**Example:**

```
host_list_ports host=tlib_host_pro125_fc0
```

**Output:**

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

| Field ID | Field output | Default position |
|---|---|---|
| **host** | Host | 1 |
| **type** | Type | 2 |
| **port_name** | Port Name | 3 |
| **num_of_visible_targets** | iSCSI targets limit | 4 |
| **num_of_visible_targets_per_network** | iSCSI targets limit per network | 5 |

### Access control

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

| User Category | Permission |
|---|---|
| Technicians | Disallowed |

## Removing a port from a host

Use the **host_remove_port** command to remove a port from a host.

```
host_remove_port host=HostName < fcaddress=wwpn | iscsi_name=iSCSIName >
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **host** | Object name | The host name. | Y |
| **fcaddress** | N/A | FC address of the port to be removed. | N |
| **iscsi_name** | iSCSI initiator name | iSCSI initiator name of the port to be removed. | N |

**Example:**

```
host_remove_port host=host1 iscsi_name=iscsi1
```

**Output:**

```
Command completed successfully
```

### Access control

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

### Return codes

- **HOST_BAD_NAME**

  The host name does not exist.
- **PORT_DOES_NOT_BELONG_TO_HOST**

  The port ID belongs to another host.
- **HOST_PORT_DOES_NOT_EXIST**

  The port ID is not defined.
- **ISCSI_HOST_ILLEGAL_PORT_NAME**

  The port name for iSCSI Host is illegal.

  **Troubleshooting:** Port names for iSCSI Hosts must contain only printable characters.
- **OLVM_LINK_IS_NOT_UP**

The IBM Hyper-Scale Mobility link is not up. The mapping list cannot be updated.

- **TARGET_NOT_CONNECTED**

  There is currently no connection to the target system.
- **TARGET_PORT_BAD_ADDRESS**

  The remote port address is illegal or does not belong to the remote target.
- **HOST_PORT_EXISTS**

  A host with this port ID is already defined.
- **MAX_PORTS_REACHED**

  The maximum number of ports defined in the system is already reached.
- **PORT_EXISTS**

  The port is already defined.
- **REMOTE_MAX_VIRTUAL_HOSTS_REACHED**

  The maximum number of defined remote virtual hosts is already reached.

# Renaming a host

Use the **host_rename** command to rename a host.

```
host_rename host=HostName new_name=Name
```

## Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **host** | Object name | The original host name. | Y |
| **new_name** | Object name | The new host name. Must be unique in the system. | Y |

The new host name must be unique in the system.

The command still succeeds even if the new name is identical to the current name.

**Example:**

```
host_rename host=server2 new_name=mailserver
```

**Output:**

```
Command completed successfully
```

## Access control

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

### Return codes

- **HOST_BAD_NAME**

  The host name does not exist.
- **HOST_NAME_EXISTS**

  The host name already exists.

## Updating a host definition

Use the **host_update** command to update a host definition.

```
host_update host=HostName [ iscsi_chap_name=iscsiChapName ] [ iscsi_chap_secret=iscsiPass ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **host** | Object name | Name that represents the host to the storage system. | Y | N/A |
| **iscsi_chap_name** | String | The host's CHAP name identifier | N | [unchanged] |
| **iscsi_chap_secret** | String | The password of the initiator used to authenticate to the storage system when CHAP is enabled | N | [unchanged] |

The command carries out the following CHAP-related checks:

- The parameters **iscsi_chap_name** and **iscsi_chap_secret** must be either both specified or both unspecified.

  These parameters have to be unique. In case they are not, an error message is displayed, but the command completes.
- The secret needs to be between 96 bits and 128 bits. There are 3 ways to enter the secret:
  - *Base64*: Requires to prefix the entry with 0b. Each subsequent character entered is treated as a 6-bit equivalent length
  - *Hex*: Requires to prefix the entry with 0x. Each subsequent character entered is treated as a 4-bit equivalent length
  - *String*: Requires no prefix (cannot be prefixed with 0b or 0x). Each character entered is treated as an 8-bit equivalent length
- If **iscsi_chap_secret** does not conform with the required secret length (96-128 bits), the command fails.

Changing the **iscsi_chap_name** and/or **iscsi_chap_secret**:

- A warning message will be displayed stating that the changes will apply only next time the host is connected.

**Example:**

```
host_update host iscsi_chap_name iscsi_chap_secret
```

**Output:**

```
Command completed successfully.
```

### Access control

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

### Warnings

- **ISCSI_CHAP_NAME_AND_SECRET_NOT_UNIQUE**

  Both iSCSI CHAP name and secret are already used by another host. Are you sure you want to reuse those values?
- **ISCSI_CHAP_SECRET_NOT_UNIQUE**

  iSCSI CHAP secret is already used by another host. Are you sure you want to reuse this value?
- **ISCSI_CHAP_NAME_NOT_UNIQUE**

  iSCSI CHAP name is already used by another host. Are you sure you want to reuse this value?

### Return codes

- **HOST_BAD_NAME**

  The host name does not exist.
- **ISCSI_CHAP_NAME_EMPTY**

  CHAP name should be a non-empty string.
- **ISCSI_CHAP_NAME_TOO_LONG**

  CHAP name is too long.
- **ISCSI_CHAP_SECRET_EMPTY**

  CHAP secret should be a non-empty string.
- **ISCSI_CHAP_SECRET_BAD_SIZE**

  CHAP secret should be 12 to 16 bytes long.
- **ISCSI_CHAP_SECRET_BAD_HEX_FORMAT**

  CHAP secret is an illegal hexadecimal number or its size is illegal. It should be 24 to 32 hexadecimal digits.

## Mapping a volume to a host or cluster

Use the **map_vol** command to map a volume to a host or a cluster.

```
map_vol <host=HostName | cluster=ClusterName> vol=VolName lun=LUN [ override=<no|yes> ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **host** | Object name | Host name. | N | N/A |
| **cluster** | Object name | Cluster name. | N | N/A |

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **vol** | Object name | Volume name. | Y | N/A |
| **lun** | Integer | LUN identifier. | Y | N/A |
| **override** | Boolean | Override the existing mapping. | N | no |

This command maps a volume to a host or to a cluster. It maps the volume to all the hosts that are contained in the cluster.

The command fails if:

- The specified host is contained in a cluster, because in this case the mapping must be done through the cluster.
- Another volume is mapped to the same LUN for this cluster/host, and the **override** parameter is not specified.
  - If the **override** parameter is specified, the mapping is replaced. The host (or all hosts in the cluster) will see continuous mapping of volume to this LUN with a different content, and probably size.
- Mapping to a cluster, if the LUN was defined as an exception.
  - Whenever the LUN is defined as an exception, map it directly to the host.

## Access control

| User Category | Permission | Condition |
|---|---|---|
| Storage administrator | Allowed | N/A |
| Storage integration administrator | Allowed | N/A |
| Application administrator | Conditionally Allowed | This volume is a snapshot. The master volume of this snapshot is mapped to a host or cluster that is associated with the user executing this command. This snapshot was created by an application administrator. |
| Security administrator | Disallowed | N/A |
| Read-only users | Disallowed | N/A |
| Technicians | Disallowed | N/A |

## Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_PERFORM_HOST_SPECIFIC_MAPPING**

  '*Host*' is part of a cluster. Are you sure you want to map this volume only for that specific host?
- **ARE_YOU_SURE_YOU_WANT_TO_MAP_VOLUME**

  Are you sure you want to map volume *Volume*, which is already mapped to another host/cluster?

## Return codes

- **HOST_BAD_NAME**

  The host name does not exist.
- **HOST_BELONGS_TO_CLUSTER**

  This host already belongs to a cluster.
- **CLUSTER_BAD_NAME**

  The cluster name does not exist.

- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **SNAPSHOT_IS_INTERNAL**

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

  Mapping conflict: the volume is already assigned.
- **LUN_ALREADY_IN_USE**

  Mapping conflict: LUN is already in use.
- **EXT_LUN_ILLEGAL**

  The LUN is out of range or does not exist.
- **VOLUME_HAS_HOST_SPECIFIC_MAPPING**

  The specified volume is currently mapped to another LUN in a host-specific mapping.
- **LUN_HAS_HOST_SPECIFIC_MAPPING**

  The specified LUN currently has another volume mapped in a host-specific mapping.
- **VOLUME_IS_NON_PROXY_OLVM_DESTINATION**

  The volume is in an IBM Hyper-Scale Mobility migration state.
- **ISCSI_HOST_ILLEGAL_PORT_NAME**

  The port name for iSCSI Host is illegal.

  **Troubleshooting:** Port names for iSCSI Hosts must contain only printable characters.
- **MAX_PORTS_REACHED**

  The maximum number of ports defined in the system is already reached.
- **OLVM_LINK_IS_NOT_UP**

  The IBM Hyper-Scale Mobility link is not up. The mapping list cannot be updated.
- **HOST_PORT_EXISTS**

  A host with this port ID is already defined.
- **OPERATION_DENIED_OBJECT_MANAGED**

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

  The maximum number of defined remote virtual hosts is already reached.
- **VOLUME_HAS_INACTIVE_DATA_MIGRATION**

  Cannot map a volume that has an inactive data migration.
- **TARGET_NOT_CONNECTED**

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

  There is currently no connection from the target system.
- **VOLUME_IS_AN_UNAVAILABLE_HYPERSWAP_PEER**

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

  Cannot associate a HyperSwap volume with a host of unconfigured type. IMPORTANT: Read the HyperSwap chapter in the 'Best Practices' document to understand the solution requirements.

# Listing the mapping of volumes to hosts, clusters, or domains

Use the **mapping_list** command to list the mapping of volumes to a specified host, cluster, or domain.

```
mapping_list [ host=HostName | cluster=ClusterName ] [ domain=DomainList ]
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **host** | Object name | List of hosts to show the mapping from. To define more than one host, separate each one with a comma. To specify all existing non-internal hosts, use "*". If no host or cluster is defined, the result returns all non-internal mappings of both hosts and clusters.<br>**Note:** Only a host or cluster argument may be provided at a time. | N | N/A |
| **cluster** | Object name | List of clusters to show the mapping from. To define more than one cluster, separate each one with a comma. To specify all existing non-internal clusters, use "*". If no host or cluster is defined, the result returns all non-internal mappings of both hosts and clusters.<br>**Note:** Only a host or cluster argument may be provided at a time. | N | N/A |
| **domain** | N/A | List of domains to show the mapping from. To define more than one domain, separate them with a comma. To specify all existing domains, use "*". | N | All user domains |

| Field ID | Field output | Default position |
|---|---|---|
| **lun** | LUN | 1 |
| **volume** | Volume | 2 |
| **proxy** | Proxy | 3 |
| **size** | Size | 4 |
| **master** | Master | 5 |
| **wwn** | WWN | 6 |
| **locked** | Locked | 7 |
| **host** | Host | 8 |

**Example:**

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

**Output:**

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

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

### Access control

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

### Return codes

- **HOST_BAD_NAME**

  The host name does not exist.
- **CLUSTER_BAD_NAME**

  The cluster name does not exist.
- **TOO_MANY_MAPPINGS**

  There are too many mappings to display.
- **DOMAIN_DOESNT_EXIST**

  The domain does not exist.

## Setting the special type of hosts or clusters

Use the **special_type_set** command to set the special type of a host or a cluster.

```
special_type_set <host=HostName | cluster=ClusterName>  type=<default|hpux|zvm|Linux|ESXi|
Windows2008|AllOther
s>
```

## Parameters

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

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

**Note:**

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

**Note:**

Modifying the **type** parameter to 'Linux' or 'ESXi' is only permitted if host or cluster are not mapped with a LUN greater than 255.

### Example:

```
special_type_set host=tlib_host_pro26_fc0 type=zvm
```

### Output:

```
Command completed successfully.
```

## Access control

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

## Return codes

- **HOST_BAD_NAME**

  The host name does not exist.
- **HOST_BELONGS_TO_CLUSTER**

  This host already belongs to a cluster.
- **HOST_HAS_MAPPED_VOLUMES**

  The host already has mapped volumes.
- **CLUSTER_BAD_NAME**

  The cluster name does not exist.
- **CLUSTER_HAS_MAPPED_VOLUMES**

  The cluster already has mapped volumes.
- **HOST_HAS_LUN_GREATER_THAN_255**

The host cannot be mapped to a LUN greater than 255 when setting type to 'Linux' or 'ESXi'.

- **CLUSTER_HAS_LUN_GREATER_THAN_255**

  The cluster cannot be mapped to a LUN greater than 255 when setting type to 'Linux' or 'ESXi'.

# Listing hosts/clusters to which a volume is mapped

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

```
vol_mapping_list vol=VolName
```

## Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **vol** | Object name | Volume name. | Y |

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

| Field ID | Field output | Default position |
|---|---|---|
| **host** | Host/Cluster | 1 |
| **type** | Type | 2 |
| **lun** | LUN | 3 |

## Access control

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

## Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.

# Unmapping a volume from a host or cluster

Use the **unmap_vol** command to unmap a volume from a host or a cluster.

```
unmap_vol <host=HostName | cluster=ClusterName> vol=VolName [ idle_seconds=IdleSeconds ]
[ force=<yes|no> ]
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **host** | Object name | Host name. | N | N/A |

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **cluster** | Object name | Cluster name. | N | N/A |
| **vol** | Object name | Volume name. | Y | N/A |
| **idle_seconds** | Integer | How many seconds the volume needs to be idle before unmapping. | N | -1 |
| **force** | Boolean | Force completing the unmap operation, even if there exists in-flight I/O that has not completed on the proxy volume. | N | no |

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

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

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

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

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

## Access control

| User Category | Permission | Condition |
|---|---|---|
| Storage administrator | Allowed | N/A |
| Storage integration administrator | Allowed | N/A |
| Application administrator | Conditionally Allowed | The volume is a snapshot, where its master volume is mapped to a host or cluster associated with the user and the snapshot was created by an application administrator. |
| Security administrator | Disallowed | N/A |
| Read-only users | Disallowed | N/A |
| Technicians | Disallowed | N/A |

## Warnings

- **HA_SLAVE_NOT_CONNECTED**

  The secondary volume in this HyperSwap relation is not connected to the primary volume. Are you sure you want to unmap the volume?

  **Troubleshooting:** It is not recommended to unmap the secondary volume if the relation is still in use.

## Return codes

- **HOST_BAD_NAME**

  The host name does not exist.
- **HOST_BELONGS_TO_CLUSTER**

  This host already belongs to a cluster.

- **CLUSTER_BAD_NAME**

  The cluster name does not exist.
- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **SNAPSHOT_IS_INTERNAL**

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

  The volume is in an IBM Hyper-Scale Mobility migration state.
- **OPERATION_DENIED_OBJECT_MANAGED**

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

  The IBM Hyper-Scale Mobility link is not up. The mapping list cannot be updated.
- **ISCSI_HOST_ILLEGAL_PORT_NAME**

  The port name for iSCSI Host is illegal.

  **Troubleshooting:** Port names for iSCSI Hosts must contain only printable characters.
- **MAX_PORTS_REACHED**

  The maximum number of ports defined in the system is already reached.
- **HOST_PORT_EXISTS**

  A host with this port ID is already defined.
- **VOLUME_NOT_IDLE**

  The volume was not idle before unmapping. Check connected hosts and idle timeout.
- **MAPPING_IS_NOT_DEFINED**

  The requested mapping is not defined.
- **REMOTE_MAX_VIRTUAL_HOSTS_REACHED**

  The maximum number of defined remote virtual hosts is already reached.
- **TARGET_NOT_CONNECTED**

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

  There is currently no connection from the target system.

# Set the default idle time before unmapping a volume

Use the **unmap_vol_set_default_idle_time** command to set the default idle time required for a volume before unmapping it.

```
unmap_vol_set_default_idle_time idle_time_seconds=IdleSeconds
```

## Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **idle_time_seconds** | Integer | Defines how many seconds the volume needs to be idle before unmapping. | Y |

**Example:**

```
unmap_vol_set_default_idle_time idle_time_seconds=10
```

**Output:**

```
Command completed successfully
```

### Access control

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

## Retrieving the default idle time before unmapping a volume

Use the **unmap_vol_get_default_idle_time** command to retrieve the default idle time required for a volume before unmapping it.

```
unmap_vol_get_default_idle_time
```

**Example:**

```
unmap_vol_get_default_idle_time
```

**Output:**

```
idle_time_seconds = "0"
```

### Access control

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

## Creating a performance class

Use the **perf_class_create** command to create a performance class.

```
perf_class_create perf_class=perfClassName [ type=<shared|independent> ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **perf_class** | String | The name of a performance class. | Y | N/A |
| **type** | Enumeration | Determines if associated objects will be limited independently or share the same limit. | N | shared |

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

**Example:**

```
perf_class_create perf_class=p1
```

**Output:**

```
Command completed successfully.
```

### Access control

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

### Return codes

- **PERF_CLASS_EXISTS**

  The performance class already exists.
- **MAX_PERF_CLASSES_REACHED**

  The maximum number of defined performance classes is already reached.

## Deleting a performance class

Use the **perf_class_delete** command to delete a performance class.

```
perf_class_delete perf_class=perfClassName
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **perf_class** | Object name | Name of a performance class. | Y |

**Example:**

```
perf_class_delete perf_class=p1
```

**Output:**

```
Command completed successfully.
```

### Access control

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

### Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_DELETE_A_PERF_CLASS**

  Are you sure you want to delete performance class *Performance Class*?
- **ARE_YOU_SURE_YOU_WANT_TO_DELETE_A_PERF_CLASS_ASSOCIATED_WITH_HOSTS**

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

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

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

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

### Return codes

- **PERF_CLASS_BAD_NAME**

  The performance class does not exist.

## Renaming a performance class

Use the **perf_class_rename** command to rename a performance class.

```
perf_class_rename perf_class=perfClassName new_name=Name
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **perf_class** | Object name | The name of an existing performance class. | Y |
| **new_name** | String | The new name for the performance class. The class new name must be unique. | Y |

**Example:**

```
perf_class_rename perf_class=p1 new_name=perf1
```

**Output:**

```
Command completed successfully.
```

### Access control

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

### Return codes

- **PERF_CLASS_EXISTS**

  The performance class already exists.
- **PERF_CLASS_BAD_NAME**

  The performance class does not exist.

## Listing details of performance classes

Use the **perf_class_list** command to list performance classes.

```
perf_class_list [ perf_class=perfClassName ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **perf_class** | String | Name of a performance class. If left unspecified, all performance classes will be listed. | N | All performance classes. |

| Field ID | Field output | Default position |
|---|---|---|
| **name** | Performance class | 1 |
| **type** | Class type | 2 |

| Field ID | Field output | Default position |
|---|---|---|
| **max_iops** | Max IO rate(IOPS) | 3 |
| **max_bw** | Max BW rate(MB/sec) | 4 |

**Example:**

```
perf_class_list
```

**Output:**

```
Performance class   Max IO rate(IOPS)   Max BW rate(MB/sec)
perf1               0                    0
```

### Access control

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

## Adding a host to a performance class

Use the **perf_class_add_host** command to add a host to a performance class.

```
perf_class_add_host perf_class=perfClassName host=HostName
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **perf_class** | Object name | The name of a performance class. | Y |
| **host** | Object name | The name of the host to be added to the performance class. | Y |

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

**Example:**

```
perf_class_add_host perf_class=p1 host=h1
```

**Output:**

```
Command completed successfully.
```

### Access control

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

### Return codes

- **PERF_CLASS_BAD_NAME**

  The performance class does not exist.
- **PERF_CLASS_ASSOCIATED_WITH_POOLS_OR_DOMAINS**

  Performance class *Performance Class* is already being used by a pool or domain.
- **PERF_CLASS_ASSOCIATED_WITH_VOLUMES**

  Performance class *Performance Class* is already being used by a volume.
- **HOST_BAD_NAME**

  The host name does not exist.
- **HOST_ALREADY_IN_PERF_CLASS**

  Host *host* is already in performance class *Performance Class*.

## Removing a host from its performance class

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

```
perf_class_remove_host host=HostName
```

### Parameters

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

**Example:**

```
perf_class_remove_host host=h1
```

**Output:**

```
Command completed successfully.
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Allowed |

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

### Return codes

- **HOST_BAD_NAME**

  The host name does not exist.
- **PERF_CLASS_DOES_NOT_CONTAIN_ANY_HOSTS**

  The performance class is already empty.

## Adding a pool to a performance class

Use the **perf_class_add_pool** command to add a pool to a performance class.

```
perf_class_add_pool perf_class=perfClassName pool=PoolName
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **perf_class** | Object name | Name of a performance class | Y |
| **pool** | Object name | Name of a pool that will be added to the performance class | Y |

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

**Example:**

```
perf_class_add_pool perf_class=p1 pool=h1
```

**Output:**

```
Command completed successfully.
```

### Access control

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

### Return codes

- **PERF_CLASS_BAD_NAME**

The performance class does not exist.

- **PERF_CLASS_ASSOCIATED_WITH_HOSTS**

  Performance class *Performance Class* is already being used by a host.
- **PERF_CLASS_ASSOCIATED_WITH_VOLUMES**

  Performance class *Performance Class* is already being used by a volume.
- **POOL_DOES_NOT_EXIST**

  The storage pool does not exist.
- **POOL_ALREADY_IN_PERF_CLASS**

  Pool *pool name* is already in performance class *Performance Class*.

## Removing a pool from its performance class

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

```
perf_class_remove_pool pool=PoolName
```

### Parameters

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

**Example:**

```
perf_class_remove_pool pool=h1
```

**Output:**

```
Command completed successfully.
```

### Access control

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

### Return codes

- **POOL_DOES_NOT_EXIST**

  The storage pool does not exist.
- **POOL_NOT_CONNECTED_TO_ANY_PERF_CLASS**

  The pool is not connected to any performance class.

# Adding a volume to a performance class

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

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

## Parameters

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

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

**Example:**

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

**Output:**

```
Command completed successfully.
```

## Access control

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

## Return codes

- **PERF_CLASS_BAD_NAME**

  The performance class does not exist.
- **PERF_CLASS_ASSOCIATED_WITH_HOSTS**

  Performance class *Performance Class* is already being used by a host.
- **PERF_CLASS_ASSOCIATED_WITH_POOLS_OR_DOMAINS**

  Performance class *Performance Class* is already being used by a pool or domain.
- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **VOLUME_ALREADY_IN_PERF_CLASS**

  Volume *volume name* is already in performance class *Performance Class*.

# Removing a volume from its performance class

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

```
perf_class_remove_vol vol=VolName
```

## Parameters

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

**Example:**

```
perf_class_remove_vol volume=h1
```

**Output:**

```
Command completed successfully.
```

## Access control

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

## Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **VOLUME_NOT_CONNECTED_TO_ANY_PERF_CLASS**

  The volume is not connected to any performance class.

# Adding a domain to a performance class

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

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

## Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **domain** | Object name | The name of the domain to be added to the performance class. | Y |

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **perf_class** | Object name | The name of a performance class. | Y |

**Example:**

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

**Output:**

```
Command completed successfully.
```

### Access control

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

### Return codes

- **DOMAIN_DOESNT_EXIST**

  The domain does not exist.
- **PERF_CLASS_BAD_NAME**

  The performance class does not exist.
- **PERF_CLASS_ASSOCIATED_WITH_HOSTS**

  Performance class *Performance Class* is already being used by a host.
- **PERF_CLASS_ASSOCIATED_WITH_VOLUMES**

  Performance class *Performance Class* is already being used by a volume.
- **DOMAIN_ALREADY_IN_PERF_CLASS**

  Domain *domain name* is already in performance class *Performance Class*.

## Removing a domain from its performance class

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

```
perf_class_remove_domain domain=DomainName
```

### Parameters

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

**Example:**

```
perf_class_remove_domain domain=d1
```

**Output:**

```
Command completed successfully.
```

### Access control

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

### Return codes

- **DOMAIN_DOESNT_EXIST**

  The domain does not exist.
- **PERF_CLASS_NOT_ASSOC_WITH_DOMAIN**

  The domain is not in any performance class.

## Setting the rate for a performance class

Use the **perf_class_set_rate** command to set the rate for a performance class.

```
perf_class_set_rate perf_class=perfClassName [ max_io_rate=iops ] [ max_bw_rate=bw ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **perf_class** | Object name | Name of a performance class. | Y | N/A |
| **max_io_rate** | Positive integer | Specifies the performance class maximum rate in IOPS per interface module (IOPS). The maximum setting allowed is 100,000. If zero is specified, the IOPS rate will not be limited. | N | Keep unchanged. |
| **max_bw_rate** | Positive integer | Specifies the performance class maximum rate per interface module (MB/sec). The maximum setting allowed is 10,000. If zero is specified, the bandwidth rate will not be limited. | N | Keep unchanged. |

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

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

**Example:**

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

**Output:**

```
Command completed successfully.
```

### Access control

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

### Return codes

- **PERF_CLASS_BAD_NAME**

  The performance class does not exist.
- **PERF_CLASS_INVALID_RATE**

  The rate set for the performance class is invalid.

## Listing host profiles

Use the **host_profile_list** command to list all host profiles.

```
host_profile_list [ host=HostName ] [ domain=DomainName ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **host** | Object name | The name of the specific host whose profiles should be listed | N | >All Host Profiles. |
| **domain** | Object name | The domain name. | N | All Domains |

The command lists all host profiles or a specific one.

| Field ID | Field output | Default position |
|---|---|---|
| **host_name** | Host Name | 1 |
| **update_time** | Update Time | 2 |
| **profile** | Profile | 3 |

**Example:**

```
host_profile_list host
```

**Output:**

```
Host Name                    Update Time          Profile
---------------------------- -------------------- ----------------------
host1                        2012-05-09 22:54:36  Windows 7
```

### Access control

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

## Updating the host profile

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

```
host_profile_set profile_value=Profile
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **profile_value** | String | The host profile value length up to 1024 characters | Y |

**Example:**

```
host_profile_set profile_value=Profile
```

**Output:**

```
Command completed successfully.
```

### Access control

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

### Return codes

- **HOST_PROFILE_UPDATE_TOO_FREQUENT**

  Host Profile has been set too often.

  **Troubleshooting:** Try again after the minimal update interval time.
- **HOST_BAD_NAME**

  The host name does not exist.
- **MAX_HOST_PROFILES_REACHED**

  The maximum number of defined host profiles is already reached.

## Removing the profile of the specified host

Use the **host_profile_clear** command to remove the profile of the specified host.

```
host_profile_clear  host=HostName
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **host** | Object name | The host name. | Y |

**Example:**

```
host_profile_clear host
```

### Access control

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

### Return codes

- **HOST_BAD_NAME**

  The host name does not exist.
- **HOST_PROFILE_DOES_NOT_EXIST**

  No profile is defined for the requested host.

## Enabling the host profiler

Use the **host_profiler_enable** command to enable the host profiler functionality.

```
host_profiler_enable
```

**Example:**

```
host_profiler_enable
```

### Access control

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

## Disabling the host profiler

Use the **host_profiler_disable** command to disable the host profiler functionality.

```
host_profiler_disable
```

**Example:**

```
host_profiler_disable
```

### Access control

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

# 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**

  The volume name does not exist.
- **LOCAL_PEER_IS_NOT_MASTER**

  The local peer is not primary.

## Listing reservation keys

Use the **reservation_key_list** command to list reservation keys.

```
reservation_key_list  [ vol=VolName ]
```

© Copyright IBM Corp. 2016, 2022

### 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**

  The volume name does not exist.

## Listing volume reservations

Use the **reservation_list** command to list volume reservations.

```
reservation_list  [ vol=VolName ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **vol** | Object name | The name of the volume whose reservations are to be listed. | N | All volumes. |

**Example:**

```
reservation_list vol=Vol1
```

**Output:**

```
Volume Name   Reserving Port   Reservation Type   Persistent
vol1                           none               none

cont:
Reservation Type   Persistent Access Type   Initiator UID   PR Generation
                   none                    -1              0
```

| Field ID | Field output | Description | Default position |
|---|---|---|---|
| **name** | Volume Name | N/A | 1 |
| **reserved_by_port** | Reserving Port | N/A | 2 |
| **reserved_by_port_isid** | Reserving ISID | N/A | 3 |
| **reservation_type** | Reservation Type | N/A | 4 |
| **persistent_reservation_type** | Persistent Reservation Type | N/A | 5 |
| **access_type** | Persistent Access Type | N/A | 6 |
| **reserving_initiator_uid** | Initiator UID | uid of reserving host | 7 |
| **pr_generation** | PR Generation | N/A | 8 |
| **reservation_age** | Reservation Age | N/A | 9 |

### Access control

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

### Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.

## Copying volumes

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

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

## Parameters

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

This command copies a source volume onto a target volume.

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

This command performs the following as a single atomic action:

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

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

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

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

The command fails in the following cases:

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

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

**Example:**

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

**Output:**

```
Command completed successfully.
```

## Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |

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

## Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_COPY_VOLUME**

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

## Return codes

- **NOT_ENOUGH_SPACE**

  No space to allocate for the volume's current usage.
- **SOURCE_VOLUME_BAD_NAME**

  The source volume name does not exist.
- **SOURCE_VOLUME_DATA_MIGRATION_UNSYNCHRONIZED**

  Data Migration to source volume has not completed.
- **TARGET_VOLUME_BAD_NAME**

  The target volume name does not exist.
- **TARGET_VOLUME_LOCKED**

  The target volume is locked.
- **TARGET_VOLUME_HAS_MIRROR**

  A mirror is defined for the target volume.
- **TARGET_VOLUME_HAS_DATA_MIGRATION**

  Data Migration is defined for the target volume.
- **VOLUME_IS_SNAPSHOT**

  The operation is not permitted on snapshots.
- **VOLUME_IDENTICAL**

  The same volume is defined as source and target.
- **VOLUME_HAS_SNAPSHOTS**

  The volume has snapshots.
- **VOLUME_IS_NOT_CONSISTENT_SLAVE**

  The operation not allowed on an inconsistent secondary volume.
- **VOLUME_IS_NOT_CONSISTENT_OLVM_DESTINATION**

  The operation not allowed on an inconsistent IBM Hyper-Scale Mobility volume.
- **TARGET_VOLUME_NOT_FORMATTED**

  The target volume is not formatted.
- **SNAPSHOT_IS_FORMATTED**

  The snapshot is formatted.
- **VOLUME_TOO_BIG_TO_COPY**

The volume is too large to be copied.

- **TARGET_VOLUME_HAS_OLVM**

  This target volume is part of an IBM Hyper-Scale Mobility relationship.
- **VOLUME_IS_OLVM_PROXY**

  The volume is in an IBM Hyper-Scale Mobility Proxy phase.
- **OPERATION_DENIED_OBJECT_MANAGED**

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

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support
- **SYSTEM_OUT_OF_PHYSICAL_SPACE**

  The operation not allowed while the system is out of physical space.
- **VOLUME_TOO_BIG**

  No space to allocate to the volume.
- **VOLUME_HAS_HA**

  This operation is forbidden on a volume with a HyperSwap relation.
- **TARGET_VOLUME_HAS_HA**

  This operation is forbidden, if the target volume is a peer in a HyperSwap relation.

# Creating a volume

Use the **vol_create** command to create a new volume.

```
vol_create vol=VolName < size=GB | size_blocks=BLOCKS > pool=PoolName  [ ext_id=Identifier ]
 [ perf_class=perfClassName ]
```

## Parameters

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

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

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

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

Volumes are created in increments of approximately 1 GB. In some cases, rounding of up to 5% of the total volume size can take place in order to improve internal accounting. The volume size is the actual "net" storage space, as seen by the user's applications, not including any internal overhead, such as rounding.

The volume is logically formatted at the creation time, which means that any read operation results in returning all zeros as a response.

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

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

**Example:**

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

**Output:**

```
Command completed successfully.
```

## Access control

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

## Warnings

- **VOLUME_SIZE_VERY_LARGE_ARE_YOU_SURE**

  The volume size is very large. It may not be possible to mirror this volume to older versions of the storage system. Are you sure?

## Return codes

- **VOLUME_CANNOT_HAVE_ZERO_SIZE**

  The volume size cannot be zero.
- **POOL_DOES_NOT_EXIST**

  The storage pool does not exist.
- **VOLUME_EXISTS**

  The volume name already exists.
- **VOLUME_BAD_PREFIX**

  The volume name has a reserved prefix.
- **VOLUME_TOO_BIG**

  No space to allocate to the volume.

- **MAX_VOLUMES_REACHED**

  The maximum allowed number of volumes is already reached.
- **ELECTRONIC_LICENSE_NOT_APPROVED**

  Operation blocked until Electronic license approval

  **Troubleshooting:** Please retrieve Electronic license version and accept it
- **VOLUME_SIZE_ABOVE_LIMIT**

  The specified volume size is above the limit.
- **INVALID_SLICE_OFFSET**

  Slice offset is illegal.
- **OPERATION_DENIED_OBJECT_MANAGED**

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

  The system is in the process of changing the encryption activation state.
- **DOMAIN_MAX_VOLUMES_REACHED**

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

  Performance class *Performance Class* is already being used by a host.
- **PERF_CLASS_ASSOCIATED_WITH_POOLS_OR_DOMAINS**

  Performance class *Performance Class* is already being used by a pool or domain.
- **PERF_CLASS_BAD_NAME**

  The performance class does not exist.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support
- **SYSTEM_OUT_OF_PHYSICAL_SPACE**

  The operation not allowed while the system is out of physical space.

# Deleting a volume

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

```
vol_delete vol=VolName
```

### Parameters

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

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

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

All storage space allocated (or reserved) for the volume is freed and returned to the volume's storage pool.

The volume is removed from all LUN Maps that contain a mapping of the volume.

This command deletes all snapshots associated with this volume. Even snapshots that are part of a snapshot group (this can happen when the volume was in a consistency group and was removed from it prior to the deletion).

This command cannot be applied to a volume that is part of a consistency group or to a volume that is mapped to a host or cluster.

The command succeeds regardless of the volume's lock state.

**Example:**

```
vol_delete vol=DBVolumeCopy
```

**Output:**

```
Command completed successfully
```

## Access control

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

## Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_DELETE_VOLUME**

  Are you sure you want to delete volume *Volume*?
- **ARE_YOU_SURE_YOU_WANT_TO_DELETE_VOLUME_WITH_SNAPSHOTS**

  Volume *Volume* has snapshots! Are you sure you want to delete this volume AND all of its snapshots?

## Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **VOLUME_HAS_MIRROR**

  A mirror is defined for this volume.
- **SNAPSHOT_IS_PART_OF_SNAPSHOT_GROUP**

  The snapshot is part of a snapshot group.
- **SNAPSHOT_IS_INTERNAL**

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

  The volume belongs to a consistency group.
- **VOLUME_IS_MAPPED**

The volume mapped to a host cannot be deleted.

- **VOLUME_IS_BOUND**

  The volume is bound to an ALU.

  **Troubleshooting:** Unbind the volume from the ALU.
- **VOLUME_HAS_MAPPED_SNAPSHOT**

  A volume with a snapshot that is mapped to a host cannot be deleted.
- **SNAPSHOT_HAS_ACTIVE_SYNC_JOB**

  The snapshot is currently the target of an active sync job.

  **Troubleshooting:** Please wait for the sync job to complete.
- **SNAPSHOT_IS_CONSISTENT_ELCS**

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

  An IBM Hyper-Scale Mobility relationship is defined for this volume.
- **VOLUME_IS_OLVM_PROXY**

  The volume is in an IBM Hyper-Scale Mobility Proxy phase.
- **VOLUME_IS_OLVM_DESTINATION**

  The volume is defined as an IBM Hyper-Scale Mobility destination.
- **OPERATION_DENIED_OBJECT_MANAGED**

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

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

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support
- **VOLUME_HAS_HA**

  This operation is forbidden on a volume with a HyperSwap relation.

# Formatting a volume

Use the **vol_format** command to formats a volume.

```
vol_format vol=VolName
```

## Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **vol** | Object name | Name of the volume to be formatted. | Y |

A formatted volume returns zeros as a response to any read command.

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

The formatting of the volume is done logically and no data is actually written to the physical storage space allocated for the volume. This allows the command to complete instantly.

The volume's lock state must be unlocked when the command is issued.

This command fails if the volume has snapshots associated with it, or if the volume is a snapshot, or if the volume is part of any mirroring or data migration definition.

**Example:**

```
vol_format vol=DBVolume
```

**Output:**

```
Command completed successfully.
```

## Access control

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

## Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_FORMAT_VOLUME**

  Volume *Volume* may contain data. Formatting it will cause data loss. Are you sure you want to format volume *Volume*?

## Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **VOLUME_IS_OLVM_PROXY**

  The volume is in an IBM Hyper-Scale Mobility Proxy phase.
- **VOLUME_HAS_SNAPSHOTS**

  The volume has snapshots.
- **VOLUME_IS_SNAPSHOT**

  The operation is not permitted on snapshots.
- **VOLUME_LOCKED**

  The volume is locked.
- **VOLUME_HAS_MIRROR**

  A mirror is defined for this volume.
- **VOLUME_HAS_OLVM**

  An IBM Hyper-Scale Mobility relationship is defined for this volume.
- **VOLUME_HAS_DATA_MIGRATION**

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

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

- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support

- **VOLUME_HAS_HA**

  This operation is forbidden on a volume with a HyperSwap relation.

# Listing volumes

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

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

[ domain=DomainName ] [ wwn=WWNString ]
```

## Parameters

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

This command lists volumes according to:

- Volume name
- Pool
- Consistency Group
- WWN

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

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

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

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

- **snapshot_format**

| Field ID | Field output | Default position |
|---|---|---|
| **name** | Name | 1 |
| **size** | Size (GB) | 2 |
| **size_MiB** | Size (MiB) | N/A |
| **vol_copy_type** | Copy type | 3 |
| **master_name** | Master Name | 4 |
| **cg_name** | Consistency Group | 5 |
| **pool_name** | Pool | 6 |
| **creator** | Creator | 7 |
| **written** | Written (GB) | 8 |
| **written_MiB** | Written (MiB) | N/A |
| **proxy** | Proxy | N/A |
| **capacity** | Capacity (blocks) | N/A |
| **modified** | Modified | N/A |
| **sg_name** | Snapshot Group Name | N/A |
| **delete_priority** | Deletion Priority | N/A |
| **locked** | Locked | N/A |
| **snapshot_time** | Snapshot Creation Time | N/A |
| **snapshot_time_on_master** | Master Copy Creation Time | N/A |
| **snapshot_internal_role** | Snapshot Internal Role | N/A |
| **snapshot_of** | Snapshot of | N/A |
| **sg_snapshot_of** | Snapshot of Snap Group | N/A |
| **wwn** | WWN | N/A |
| **mirrored** | Mirrored | N/A |
| **locked_by_pool** | Locked by Pool | N/A |
| **capacity_used_by_snapshots_MiB** | Capacity Used by Snapshots (MiB) | N/A |
| **short_lived_io** | Short Live IO | N/A |
| **enable_VAAI** | VAAI enabled | N/A |
| **user_disabled_VAAI** | VAAI disabled by user | N/A |
| **snapshot_format** | Snapshot Format | N/A |
| **unmap_support** | Unmap Support | N/A |
| **managed** | Managed | N/A |
| **marked** | Marked | N/A |
| **perf_class** | Performance Class Name | N/A |
| **thin_provisioning_savings** | Thin Provisioning Savings (%) | N/A |
| **est_compression_factor** | Est. Compression Factor | N/A |
| **unique_stored_data** | Unique Stored Data (GB) | N/A |
| **ha** | HA Relation | N/A |
| **target_port_group_id** | TPG ID | N/A |
| **target_port_group_state** | TPG State | N/A |
| **lock_modes** | Lock Modes | N/A |

| Field ID | Field output | Default position |
|---|---|---|
| **copy_master_wwn** | Copy Master | N/A |

**Example:**

```
vol_list
```

**Output:**

```
Name                              Size (GB)   Copy Type   Master Name
--------------------------------- ----------- ----------- ------------------
DBLog                             3006        None
Dev                               2010        Copy
Dev.snapshot_00001                2010        Snapshot    Dev
VM_1                              2010        Replicated

cont:
Consistency Group   Pool             Creator           Written (GB)
------------------- ---------------- ----------------- --------------
                    MainPool         admin             21
                    MainPool         admin             13
                    MainPool         admin             0
                    MainPool         admin             20
```

### Access control

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

## Listing a volume's extended attributes

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

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

### Parameters

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

This command lists extended attributes of volumes according to:

- Volume name

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

| Field ID | Field output | Default position |
|---|---|---|
| **uid** | UID | N/A |

**Example:**

```
vol_list_extended
```

**Output:**

```
Name                 WWN                                Product Serial Number
-------------------- ---------------------------------- -------------------------
DBLog                60017380000035C3000000000000000A   MN035C3000000000000000A
Dev                  60017380000035C3000000000000000B   MN035C3000000000000000B
Dev.snapshot_00001   60017380000035C3000000000000000D   MN035C3000000000000000D
Dev.snapshot_00002   60017380000035C3000000000000000E   MN035C3000000000000000E
Dev.snapshot_00003   60017380000035C3000000000000000F   MN035C3000000000000000F
Marketing            60017380000035C3000000000000000C   MN035C3000000000000000C
```

### Access control

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

### Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.

## Locking a volume

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

```
vol_lock vol=VolName
```

### Parameters

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

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

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

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

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

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

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

If applied on a volume that is part of an IBM Hyper-Scale Mobility relation, the command has to be acknowledged by both source and destination volumes. Otherwise, a return code is returned (see below).

**Example:**

```
vol_lock vol=DBVolume
```

**Output:**

```
Command completed successfully.
```

## Access control

| User Category | Permission | Condition |
|---|---|---|
| Storage administrator | Allowed | N/A |
| Storage integration administrator | Allowed | N/A |
| Application administrator | Conditionally Allowed | The volume is a snapshot, where its master volume is mapped to a host or cluster associated with the user and the snapshot was created by an application administrator. |
| Security administrator | Disallowed | N/A |
| Read-only users | Disallowed | N/A |
| Technicians | Disallowed | N/A |

## Return codes

- **SNAPSHOT_IS_INTERNAL**

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

  The volume name does not exist.
- **VOLUME_IS_OLVM_PROXY**

  The volume is in an IBM Hyper-Scale Mobility Proxy phase.
- **VOLUME_IS_SLAVE**

  The volume is defined as a secondary volume.
- **VOLUME_IS_OLVM_DESTINATION**

  The volume is defined as an IBM Hyper-Scale Mobility destination.
- **SNAPSHOT_IS_PART_OF_SNAPSHOT_GROUP**

  The snapshot is part of a snapshot group.
- **OPERATION_DENIED_OBJECT_MANAGED**

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

The data reduced tier is offline, the operation is not allowed.

**Troubleshooting:** Contact IBM Support

- **LOCAL_PEER_IS_NOT_MASTER**

  The local peer is not primary.

- **REMOTE_TARGET_NOT_CONNECTED**

  There is currently no connection from the target system.

- **TARGET_NOT_CONNECTED**

  There is currently no connection to the target system.

- **HA_IS_NOT_OPERATIONAL**

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

# Renaming a volume

Use the **vol_rename** command to rename a volume.

```
vol_rename vol=VolName new_name=Name
```

## Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **vol** | Object name | Name of the volume to be renamed. | Y |
| **new_name** | Object name | New volume name. | Y |

The new name of the volume must be unique in the system.

This command succeeds even if the new name is identical to the current name. It also succeeds regardless of the volume's lock state.

Renaming a snapshot does not change the name of its master volume. Renaming a master volume does not change the names of its associated snapshots.

If applied on a volume that is part of an IBM Hyper-Scale Mobility relation, the command has to be acknowledged by both source and destination volumes. Otherwise, a return code is returned (see below).

**Example:**

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

**Output:**

```
Command completed successfully
```

## Access control

| User Category | Permission | Condition |
|---|---|---|
| Storage administrator | Allowed | N/A |
| Storage integration administrator | Allowed | N/A |

| User Category | Permission | Condition |
|---|---|---|
| Application administrator | Conditionally Allowed | The volume is a snapshot, where its master volume is mapped to a host or cluster associated with the user and the snapshot was created by an application administrator. |
| Security administrator | Disallowed | N/A |
| Read-only users | Disallowed | N/A |
| Technicians | Disallowed | N/A |

## Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **VOLUME_EXISTS**

  The volume name already exists.
- **SNAPSHOT_IS_INTERNAL**

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

  The snapshot is part of a snapshot group.
- **VOLUME_BAD_PREFIX**

  The volume name has a reserved prefix.
- **SNAPSHOT_IS_CONSISTENT_ELCS**

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

  IBM Hyper-Scale Mobility error.
- **COMMAND_NOT_SUPPORTED_FOR_OLVM_VOLUMES**

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

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

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support
- **LOCAL_PEER_IS_NOT_MASTER**

  The local peer is not primary.
- **TARGET_NOT_CONNECTED**

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

  There is currently no connection from the target system.
- **HA_IS_NOT_OPERATIONAL**

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

The secondary volume with the indicated name already exists. The name cannot be reused.

- **REMOTE_ALU_EXISTS**

  An ALU with the indicated secondary volume name already exists on the remote machine.

## Resizing a volume

Use the **vol_resize** command to resize a volume.

```
vol_resize vol=VolName  < size=GB | size_blocks=BLOCKS > [ shrink_volume=<yes|no> ]
 [ force_on_inactive_mirror=<yes|no> ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **vol** | Object name | The name of the volume to be resized. | Y | N/A |
| **size** | N/A | The new volume size. | N | N/A |
| **size_blocks** | N/A | New size of volumes in number of blocks. | N | N/A |
| **shrink_volume** | Boolean | Must be specified as yes if the new size is smaller than the current size. | N | No |
| **force_on_inactive_mirror** | Boolean | The parameter is required for a successful resize of a volume if (1) the volume is mirrored, (2) the volume is a master, and (3) the mirror has been deactivated by the system following a previously issued resize command that failed to successfully complete due to a communication error. | N | No |

The volume can be resized in either direction. However, whenever the volume is downsized, you have to specify this with **shrink_volume=yes**.

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

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

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

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

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

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

Using the **force_on_inactive_mirror** parameter:

- This parameter forces the resizing of a mirror peer even if mirroring is inactive (this may happen when the mirroring cannot be activated due to size mismatch).

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

**Example:**

```
vol_resize -y vol=DBVolume size=2500
```

**Output:**

```
Command completed successfully.
```

## Access control

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

## Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_ENLARGE_VOLUME**

  Are you sure you want to increase the volume size?
- **ARE_YOU_SURE_YOU_WANT_TO_REDUCE_VOLUME**

  Decreasing the volume size may cause data loss. Are you sure you want to proceed?
- **VOLUME_SIZE_VERY_LARGE_ARE_YOU_SURE**

  The volume size is very large. It may not be possible to mirror this volume to older versions of the storage system. Are you sure?

## Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **VOLUME_TOO_BIG**

  No space to allocate to the volume.
- **REMOTE_VOLUME_SIZE_ABOVE_LIMIT**

  The specified volume size is above the limit of the remote machine.
- **VOLUME_LOCKED**

  The volume is locked.
- **VOLUME_HAS_DATA_MIGRATION**

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

  A mapped volume's size cannot be decreased.

- **CAN_NOT_SHRINK_VOLUME_WITH_SNAPSHOTS**

  The size of volume with snapshots cannot be decreased.
- **CAN_NOT_SHRINK_REMOTE_VOLUME_WITH_SNAPSHOTS**

  The remote volume has snapshots.
- **CAN_NOT_SHRINK_MAPPED_REMOTE_VOLUME**

  The remote volume is mapped.
- **VOLUME_IS_BOUND**

  The volume is bound to an ALU.

  **Troubleshooting:** Unbind the volume from the ALU.
- **REMOTE_VOLUME_HAS_DATA_MIGRATION**

  Data migration is already defined for the secondary volume.
- **VOLUME_CANNOT_HAVE_ZERO_SIZE**

  The volume size cannot be zero.
- **CAN_NOT_SHRINK_SNAPSHOTS**

  The size of snapshots cannot be decreased.
- **CAN_NOT_RESIZE_ASYNC_INTERVAL_VOLUMES**

  The size of volumes with asynchronous mirroring cannot be changed.
- **CAN_NOT_SHRINK_VOLUME**

  The size of volumes cannot be decreased without an explicit request.
- **MIRROR_SIZE_MISMATCH**

  The secondary and primary volume sizes are different.
- **MIRROR_POSSIBLE_SIZE_MISMATCH**

  The secondary and primary volume sizes may be different.
- **HA_POSSIBLE_SIZE_MISMATCH**

  Primary and secondary HyperSwap volume sizes may be different.
- **VOLUME_SIZE_ABOVE_LIMIT**

  The specified volume size is above the limit.
- **COMMAND_NOT_SUPPORTED_FOR_OLVM_VOLUMES**

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

  The mirror is non-operational.
- **VOLUME_IS_SLAVE**

  The volume is defined as a secondary volume.
- **MIRROR_RETRY_OPERATION**

  There is an operation in progress on this mirror.

  **Troubleshooting:** Retry the command in a few seconds.
- **VOLUME_HAS_MULTIPLE_RELATIONS**

  The volume has multiple relations. The operation is not allowed or a target must be specified.
- **REMOTE_MIRROR_IS_STANDBY**

  The remote mirror is marked as Standby.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

The data reduced tier is offline, the operation is not allowed.

**Troubleshooting:** Contact IBM Support

- **LOCAL_PEER_IS_NOT_MASTER**

  The local peer is not primary.

- **TARGET_NOT_CONNECTED**

  There is currently no connection to the target system.

- **REMOTE_TARGET_NOT_CONNECTED**

  There is currently no connection from the target system.

- **HA_IS_NOT_OPERATIONAL**

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

- **HA_RETRY_OPERATION**

  An operation is in progress on this HyperSwap relation.

  **Troubleshooting:** Try issuing the command again in a few seconds.

- **VOLUME_CANNOT_BE_RESIZED_DUE_TO_CONTAINING_CONS_GROUP_SIZE_LIMIT**

  The requested volume size will cause the containing CG to exceed the maximal CG size.

# Unlocking a volume

Use the **vol_unlock** command to unlock a volume, so that it is no longer read-only and can be written to.

```
vol_unlock vol=VolName
```

## Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **vol** | Object name | The name of the volume to unlock. | Y |

An unlocked volume is no longer write-protected.

The lock state of regular volumes is set to *unlocked* when they are created. The lock state of snapshots is set to *locked* when they are created.

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

The modification time is the time when the unlock command was executed, regardless of the actual changes performed on the volume via write commands.

If applied on a volume that is part of an IBM Hyper-Scale Mobility relation, the command has to be acknowledged by both source and destination volumes. Otherwise, a return code is returned (see below).

**Example:**

```
vol_unlock vol=DBVolume
```

**Output:**

```
Command completed successfully.
```

## Access control

| User Category | Permission | Condition |
|---|---|---|
| Storage administrator | Allowed | N/A |
| Storage integration administrator | Allowed | N/A |
| Application administrator | Conditionally Allowed | The volume is a snapshot, where its master volume is mapped to a host or cluster associated with the user and the snapshot was created by an application administrator. |
| Security administrator | Disallowed | N/A |
| Read-only users | Disallowed | N/A |
| Technicians | Disallowed | N/A |

## Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_UNLOCK_SNAPSHOT**

  Are you sure you want to unlock snapshot *Snapshot*?

## Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **VOLUME_IS_OLVM_PROXY**

  The volume is in an IBM Hyper-Scale Mobility Proxy phase.
- **VOLUME_IS_SLAVE**

  The volume is defined as a secondary volume.
- **VOLUME_IS_OLVM_DESTINATION**

  The volume is defined as an IBM Hyper-Scale Mobility destination.
- **SNAPSHOT_IS_INTERNAL**

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

  The snapshot is part of a snapshot group.
- **OPERATION_DENIED_OBJECT_MANAGED**

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

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support
- **SYSTEM_OUT_OF_PHYSICAL_SPACE**

  The operation not allowed while the system is out of physical space.
- **LOCAL_PEER_IS_NOT_MASTER**

  The local peer is not primary.

- **TARGET_NOT_CONNECTED**

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

  There is currently no connection from the target system.
- **HA_IS_NOT_OPERATIONAL**

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

# 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 |
| 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 |

© Copyright IBM Corp. 2016, 2022

| User Category | Permission | Condition |
|---|---|---|
| Read-only users | Disallowed | N/A |
| Technicians | Disallowed | N/A |

### Return codes

- **VOLUME_IS_NOT_A_SNAPSHOT**

  The operation is permitted on snapshots only.
- **SNAPSHOT_ILLEGAL_PRIORITY**

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

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

  The snapshot is part of a snapshot group.
- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **SNAPSHOT_IS_CONSISTENT_ELCS**

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

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

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support

## Creating a snapshot

Use the **snapshot_create** command to create a snapshot of an existing volume.

```
snapshot_create vol=VolName  < [ name=Name ] [ delete_priority=del_value ]
 > | < overwrite=Name > [ ext_id=Identifier ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **vol** | Object name | Name of the volumes to snapshot. | Y | N/A |
| **name** | Object name | Names of the new snapshots. | N | Auto-generated names. |
| **delete_priority** | Integer | The deletion priority of the volume's snapshot. | N | 1 |
| **overwrite** | Object name | Name of an existing snapshot to be overwritten with the current volumes content. | N | N/A |
| **ext_id** | String | External identifier of the volume. | N | N/A |

This command creates a new snapshot for an existing volume, which is referred to as the snapshot's master volume. The snapshot's content is the same as the master volume at the exact point in time when the snapshot was created. The snapshot remains unchanged, although the master volume keeps changing after the snapshot is created. Upon a successful completion of this command, the snapshot is created and assigned a name that can later be used by other commands. The name does not have to be new. It can be the name of an already existing snapshot (in such a case, the already existing snapshot is overridden).

A write operation can be processed at the exact time of the snapshot creation, meaning that the write operation request was sent to the system before the command was executed, while the write was acknowledged after the command was executed. In this case, the content of the snapshot is not deterministic and may either contain the original value before the write operation, or the new value after the write operation. In fact, the snapshot's data may even contain a mixture of the two, where some blocks are equal to the volume before the write operation and other blocks are equal to the value after the write operation.

The new snapshot is initially locked for changes.

The created snapshot acts like a regular volume, except for the differences described below:

- The snapshot's name is either automatically generated from its master volume's name or given as a parameter to the command. It can later be changed without altering the snapshot's modification state.
- Upon successful completion of the command, the system assigns a unique SCSI ID to the snapshot. The creation time of the snapshot is set to the current time and is never changed until the snapshot is deleted.
- The size of the snapshot is the same as its master volume's size, but no storage space is reserved for the snapshot. This means that the functionality of the snapshot is not guaranteed. When the snapshot's storage pool is exhausted, the snapshot may be deleted.
- The snapshot's lock state is initially set to "locked", and as long as it is not "unlocked", the snapshot remains an exact image of the master volume at creation time and can be the source for a restore operation. The modification state of the snapshot is initially set to "unmodified".

During creation, the snapshot's deletion priority can be set explicitly, or it is automatically set to the default value. The deletion priority determines which snapshots will be deleted first when the storage pool runs out of snapshot storage. This may happen due to the redirect-on-write mechanisms which share unchanged data between volumes and their snapshots, as well as between snapshots of the same volume.

The Auto Delete Priority can have a value between 1 and 4, as follows:

- 1 = Is last to be deleted automatically ("1" is the default set by the system)
- ...
- 4 = Is first to be deleted automatically

The snapshot is associated with its master volume and this association cannot be broken or changed as long as the snapshot exists.

The `overwrite` option copies the current content of the volume into one of its existing snapshots (set as an input argument). The overwritten snapshot keeps the same SCSI device WWN and same mapping, so hosts maintain a continuous mapping to the snapshot, without any need for a rescan or similar operation. The overwritten snapshot must be an existing snapshot of the given volume. The overwritten snapshot cannot be part of a snapshot group.

This command fails when no snapshot space is defined in the storage pool the master volume belongs to.

Mirroring limitations:

- This command fails if the volume is a slave of an asynchronous mirroring coupling.
- This command fails if the volume is a slave of an inconsistent synchronous coupling.

**Example:**

```
snapshot_create vol=DBVolume name=DBVolume.snapshot1 delete_priority=2
```

**Output:**

```
Command completed successfully.
```

## Access control

| User Category | Permission | Condition |
|---|---|---|
| Storage administrator | Allowed | N/A |
| Storage integration administrator | Allowed | N/A |
| Application administrator | Conditionally Allowed | The volume is mapped to a host or a cluster associated with the user. If a snapshot overwrite is used, the target snapshot must be one created by a server administrator. |
| Security administrator | Disallowed | N/A |
| Read-only users | Disallowed | N/A |
| Technicians | Disallowed | N/A |

## Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **MAX_VOLUMES_REACHED**

  The maximum allowed number of volumes is already reached.
- **DOMAIN_MAX_VOLUMES_REACHED**

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

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

  The operation is not permitted on snapshots.
- **VOLUME_EXISTS**

  The volume name already exists.
- **VOLUME_BAD_PREFIX**

  The volume name has a reserved prefix.
- **VOLUME_DATA_MIGRATION_UNSYNCHRONIZED**

  Data Migration to this volume has not completed.
- **OVERWRITE_SNAPSHOT_BAD_NAME**

  The snapshot name does not exist.
- **OVERWRITE_SNAPSHOT_IS_MASTER_VOL**

  This snapshot cannot be overwritten because it is a primary volume.
- **SNAPSHOT_OVERWRITE_MISMATCH**

  The specified snapshot is not a snapshot of the specified volume.
- **SNAPSHOT_IS_PART_OF_SNAPSHOT_GROUP**

The snapshot is part of a snapshot group.

- **SNAPSHOT_IS_INTERNAL**

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

  There is not enough space to create a snapshot.
- **VOLUME_IS_NOT_CONSISTENT_SLAVE**

  The operation not allowed on an inconsistent secondary volume.
- **VOLUME_IS_NOT_CONSISTENT_OLVM_DESTINATION**

  The operation not allowed on an inconsistent IBM Hyper-Scale Mobility volume.
- **VOLUME_IS_OLVM_PROXY**

  The volume is in an IBM Hyper-Scale Mobility Proxy phase.
- **SNAPSHOT_HAS_ACTIVE_SYNC_JOB**

  The snapshot is currently the target of an active sync job.

  **Troubleshooting:** Please wait for the sync job to complete.
- **TOO_MANY_FAST_SNAPSHOTS_IN_VOLUME**

  The maximum allowed number of fast snapshots for this volume is already reached.
- **OPERATION_DENIED_OBJECT_MANAGED**

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

  Cannot create all the volumes, because otherwise the number of volumes will exceed the allowed maximum.
- **DOMAIN_WILL_EXCEED_MAXIMUM_VOLUMES_ALLOWED**

  Cannot create all the volumes, because otherwise the maximum allowed number of volumes in the domain will be exceeded.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support
- **MAX_SNAPSHOTS_PER_VOLUME_REACHED**

  The maximum allowed number of snapshots is already reached.

# Deleting a snapshot

Use the **snapshot_delete** command to delete a snapshot.

```
snapshot_delete snapshot=SnapshotName
```

## Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **snapshot** | Object name | Snapshot to be deleted. | Y |

This command cannot be used to delete a master volume, or a snapshot which is mapped to a host or cluster, or an internal snapshot of a mirroring.

**Example:**

```
snapshot_delete snapshot=DBVolume.snapshot1
```

**Output:**

```
Command completed successfully
```

## Access control

| User Category | Permission | Condition |
|---|---|---|
| Storage administrator | Allowed | N/A |
| Storage integration administrator | Allowed | N/A |
| Application administrator | Conditionally Allowed | The master volume of the snapshot is mapped to a host or cluster associated with the user and the snapshot was created by the application administrator. |
| Security administrator | Disallowed | N/A |
| Read-only users | Disallowed | N/A |
| Technicians | Disallowed | N/A |

## Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **VOLUME_IS_NOT_A_SNAPSHOT**

  The operation is permitted on snapshots only.
- **SNAPSHOT_IS_INTERNAL**

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

  The snapshot is part of a snapshot group.
- **SNAPSHOT_IS_MAPPED**

  A snapshot that is mapped to a host cannot be deleted.
- **VOLUME_IS_BOUND**

  The volume is bound to an ALU.

  **Troubleshooting:** Unbind the volume from the ALU.
- **SNAPSHOT_HAS_ACTIVE_SYNC_JOB**

  The snapshot is currently the target of an active sync job.

  **Troubleshooting:** Please wait for the sync job to complete.
- **SNAPSHOT_IS_CONSISTENT_ELCS**

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

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

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support

# Duplicating a snapshot

Use the **snapshot_duplicate** command to duplicate an existing snapshot.

```
snapshot_duplicate snapshot=SnapshotName [ name=Name ]
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **snapshot** | Object name | The name of the snapshot to duplicate. | Y | N/A |
| **name** | Object name | Name of the new snapshot to be generated. | N | Automatically generated name. |

The newly created snapshot is initially locked for changes and is associated with the master volume of the existing snapshot. The content of the newly created snapshot is identical with the content of the source snapshot.

It is useful to duplicate a snapshot before unlocking it for write operations. The duplicate snapshot can be used as a logical backup of the data in case the write operation caused logical data corruption.

Upon successful completion of the command, a new duplicate snapshot is created.

The duplicated snapshot is identical with the source snapshot. It has the same creation time and behaves as if it was created at the exact same moment and from the same master volume.

The duplicate snapshot's name is either automatically generated from its master volume's name or provided as a parameter. It can later be changed without altering its modification state.

A snapshot can be duplicated multiple times. A duplicated snapshot can be the source for further duplications.

**Example:**

```
snapshot_duplicate snapshot=DBVolume.snapshot1 name=DBVolume.snapshot1.copy
```

**Output:**

```
Command completed successfully.
```

## Access control

| User Category | Permission | Condition |
|---|---|---|
| Storage administrator | Allowed | N/A |
| Storage integration administrator | Allowed | N/A |
| Application administrator | Conditionally Allowed | The master volume of the snapshot is mapped to a host or cluster associated with the user and the snapshot was created by the application administrator. |
| Security administrator | Disallowed | N/A |
| Read-only users | Disallowed | N/A |
| Technicians | Disallowed | N/A |

## Return codes

- **VOLUME_BAD_NAME**

The volume name does not exist.

- **MAX_VOLUMES_REACHED**

  The maximum allowed number of volumes is already reached.
- **DOMAIN_MAX_VOLUMES_REACHED**

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

  The operation is permitted on snapshots only.
- **VOLUME_EXISTS**

  The volume name already exists.
- **SNAPSHOT_IS_PART_OF_SNAPSHOT_GROUP**

  The snapshot is part of a snapshot group.
- **VOLUME_BAD_PREFIX**

  The volume name has a reserved prefix.
- **OPERATION_DENIED_OBJECT_MANAGED**

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

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support
- **MAX_SNAPSHOTS_PER_VOLUME_REACHED**

  The maximum allowed number of snapshots is already reached.

## Formatting a snapshot

Use the **snapshot_format** command to format a snapshot.

```
snapshot_format snapshot=SnapshotName
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **snapshot** | Object name | The snapshot to be formatted. | Y |

This command deletes the content of a snapshot while maintaining its mapping to the host. The format operation results with:

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

**Example:**

```
snapshot_format snapshot
```

**Output:**

```
Command completed successfully.
```

### Access control

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

### Return codes

- **SNAPSHOT_IS_PART_OF_SNAPSHOT_GROUP**

  The snapshot is part of a snapshot group.
- **SNAPSHOT_HAS_ACTIVE_SYNC_JOB**

  The snapshot is currently the target of an active sync job.

  **Troubleshooting:** Please wait for the sync job to complete.
- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **SNAPSHOT_IS_INTERNAL**

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

  The maximum allowed number of volumes is already reached.
- **SNAPSHOT_IS_FORMATTED**

  The snapshot is formatted.
- **ELCS_CANNOT_BE_FORMATTED**

  The snapshot is an ELCS and cannot be formatted.
- **VOLUME_IS_NOT_A_SNAPSHOT**

  The operation is permitted on snapshots only.
- **OPERATION_DENIED_OBJECT_MANAGED**

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

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support

## Listing snapshot information

Use the **snapshot_list** command to list snapshot information.

```
snapshot_list vol=VolName [ domain=DomainName ]
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **vol** | Object name | List of all the snapshots of this volume. | Y | N/A |
| **domain** | Object name | The domain name. | N | All Domains |

This command lists snapshot information for all the snapshots of a specified volume.

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

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

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

- **snapshot_format**

| Field ID | Field output | Default position |
|---|---|---|
| **name** | Name | 1 |
| **size** | Size (GB) | 2 |
| **size_MiB** | Size (MiB) | N/A |
| **vol_copy_type** | Copy type | 3 |
| **master_name** | Master Name | 4 |
| **cg_name** | Consistency Group | 5 |
| **pool_name** | Pool | 6 |
| **creator** | Creator | 7 |
| **written** | Written (GB) | 8 |
| **written_MiB** | Written (MiB) | N/A |
| **proxy** | Proxy | N/A |
| **capacity** | Capacity (blocks) | N/A |
| **modified** | Modified | N/A |
| **sg_name** | Snapshot Group Name | N/A |
| **delete_priority** | Deletion Priority | N/A |
| **locked** | Locked | N/A |
| **snapshot_time** | Snapshot Creation Time | N/A |
| **snapshot_time_on_master** | Master Copy Creation Time | N/A |
| **snapshot_internal_role** | Snapshot Internal Role | N/A |
| **snapshot_of** | Snapshot of | N/A |
| **sg_snapshot_of** | Snapshot of Snap Group | N/A |
| **wwn** | WWN | N/A |
| **mirrored** | Mirrored | N/A |
| **locked_by_pool** | Locked by Pool | N/A |
| **capacity_used_by_snapshots_MiB** | Capacity Used by Snapshots (MiB) | N/A |
| **short_lived_io** | Short Live IO | N/A |
| **enable_VAAI** | VAAI enabled | N/A |
| **user_disabled_VAAI** | VAAI disabled by user | N/A |
| **snapshot_format** | Snapshot Format | N/A |

| Field ID | Field output | Default position |
|---|---|---|
| **unmap_support** | Unmap Support | N/A |
| **managed** | Managed | N/A |
| **marked** | Marked | N/A |
| **perf_class** | Performance Class Name | N/A |
| **thin_provisioning_savings** | Thin Provisioning Savings (%) | N/A |
| **est_compression_factor** | Est. Compression Factor | N/A |
| **unique_stored_data** | Unique Stored Data (GB) | N/A |
| **ha** | HA Relation | N/A |
| **target_port_group_id** | TPG ID | N/A |
| **target_port_group_state** | TPG State | N/A |
| **lock_modes** | Lock Modes | N/A |
| **copy_master_wwn** | Copy Master | N/A |

**Example:**

```
snapshot_list vol=DBVolume
```

**Output:**

```
Name               Size (GB)  Master Name  Consistency Group  Pool
DBVolume.sp1       2508       DBVolume                        default
DBVolume.sp1.copy  2508       DBVolume                        default
```

### Access control

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

## Restoring a volume from a snapshot

Use the **snapshot_restore** command to restore a master volume or a snapshot from one of its associated snapshots.

```
snapshot_restore snapshot=SnapshotName [ target_snapshot=SnapshotName ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **snapshot** | Object name | Name of the snapshot with which to restore its master volume, or snapshot. | Y | N/A |

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **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**

  The volume name does not exist.
- **VOLUME_IS_NOT_A_SNAPSHOT**

  The operation is permitted on snapshots only.
- **VOLUME_TOO_BIG**

  No space to allocate to the volume.
- **SNAPSHOT_IS_PART_OF_SNAPSHOT_GROUP**

  The snapshot is part of a snapshot group.
- **VOLUME_HAS_MIRROR**

  A mirror is defined for this volume.
- **VOLUME_HAS_HA**

  This operation is forbidden on a volume with a HyperSwap relation.
- **VOLUME_LOCKED**

  The volume is locked.
- **SNAPSHOTS_BELONG_TO_DIFFERENT_MASTERS**

  The target and source snapshots must be snapshots of the same volume.
- **TARGET_SNAPSHOT_BAD_NAME**

  The target snapshot name does not exist.
- **TARGET_SNAPSHOT_IS_PART_OF_SNAPSHOT_GROUP**

  The target snapshot is part of a snapshot group.
- **TARGET_SNAPSHOT_IS_MASTER**

  The target snapshot is a primary volume.
- **TARGET_SNAPSHOT_IS_OLVM_DESTINATION**

  The target snapshot is an IBM Hyper-Scale Mobility destination volume.
- **TARGET_SNAPSHOT_IS_OLVM_PROXY**

  The target snapshot is an IBM Hyper-Scale Mobility proxy volume.
- **TARGET_SNAPSHOT_SAME_AS_SNAPSHOT**

  The source snapshot must be different from the target snapshot.
- **TARGET_SNAPSHOT_HAS_ACTIVE_SYNC_JOB**

  The target snapshot is currently the target of an active sync job.

  **Troubleshooting:** Please wait for sync job to complete
- **OPERATION_DENIED_OBJECT_MANAGED**

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

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support

# 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 |

The consistency group can contain up to 128 volumes. The maximum allowed total effective capacity of the volumes in a consistency group is limited to 1 PB.

Requirements for successful command completion:

- The volume and consistency group are associated with the same pool.
- The volume is not already part of a consistency group.
- The volume is not a snapshot.
- The consistency group has less than the maximum number of volumes (see above).

Adding a mirrored volume to a non-mirrored consistency group:

- Such an addition always succeeds and the volume will retain its mirroring settings.

Requirements for successful command completion for a mirrored consistency group:

- The command must be issued only on the master consistency group.
- The command cannot be run during the initialization of the volume or consistency group.
- The volume does not have any outstanding ad-hoc sync jobs.
- The volume has to be mirrored, and its following mirroring settings must be identical to those of the consistency group: mirroring type (for example, synchronous), mirroring status, mirroring target, target pool, designation.
- In addition, for a mirrored consistency group that is defined as `sync_best_effort` (synchronous):
  - The synchronization status of both volume and consistency group has to be Synchronized.
- For a mirrored consistency group that is defined as `async_interval` (asynchronous):
  - The volume and consistency group must have the following identical settings and values: schedule, remote schedule, timestamp of the last replicated snapshot.
  - The synchronization status of the volume and consistency group must be RPO_OK

© Copyright IBM Corp. 2016, 2022

- The link has to be up.

Adding a mirrored volume to a mirrored volume and consistency group also adds the volume's peer to the volume and consistency group's peer. Once added, the mirrored volume will be set the RPO of the mirrored volume and consistency group.

The mirrored consistency group has one sync job for all pertinent mirrored volumes within the consistency group.

If the command **cg_add_vol** is issued on a mirrored master consistency group, which fails to receive an acknowledgment from the slave until the command times out or due to an unexpected failure, the **MIRROR_POSSIBLE_CONS_GROUP_MEMBERSHIP_MISMATCH** completion code is returned. The completion code indicates that the member lists of the mirror consistency group peers might not be the same.

If the volume to be added is part of an IBM Hyper-Scale Mobility relation, the command is applicable only to the destination volume. For the command to complete successfully, the destination volume must be in a proxy state.

**Note:**

In rare cases, when adding a volume to a consistency group (CG) with the **cg_add_vol** command, the operation might fail with the return code **MIRROR_LAST_SYNC_TIMES_DIFFER** recorded. To avoid this issue, retry the **cg_add_vol** command after the next volume mirror sync job completes, by which time the synchronization times should match.

**Example:**

```
cg_add_vol cg=DBGroup vol=DBLog
```

**Output:**

```
Command completed successfully.
```

## Access control

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

## Warnings

- **CONS_GROUP_MIRROR_DESIGNATION_MISMATCH**

  The volume's role in a mirroring or HyperSwap relation is different from the consistency group role. Are you sure you want to add the volume to that consistency group?

## Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **CONS_GROUP_BAD_NAME**

  The consistency group name does not exist.

- **CONS_GROUP_IS_SLAVE**

  The consistency group's role in a mirroring relationship is secondary.
- **CONS_GROUP_MIRROR_STATE_MISMATCH**

  All volumes in a mirrored consistency group must have the same mirroring sync state.
- **MAX_VOLUMES_IN_CONS_GROUP_REACHED**

  The consistency group contains the maximum allowed number of volumes.
- **MAX_VOLUMES_IN_REMOTE_CONS_GROUP_REACHED**

  The remote consistency group contains the maximum allowed number of volumes.
- **MIRROR_CONS_GROUP_CANNOT_CONTAIN_NON_MIRRORED_VOLS**

  A mirrored consistency group cannot include volumes that are not mirrored.
- **MIRROR_HAS_SYNC_JOB**

  The operation is not permitted on a mirror with active sync jobs.
- **MIRROR_IS_NOT_SYNCHRONIZED**

  The mirror is not synchronized.
- **MIRROR_LAST_SYNC_TIMES_DIFFER**

  All mirrors must have the same last sync time.
- **MIRROR_RETRY_OPERATION**

  There is an operation in progress on this mirror.

  **Troubleshooting:** Retry the command in a few seconds.
- **REMOTE_VOLUME_BAD_POOL**

  The remote volume and remote consistency group belong to different storage pools.
- **REMOTE_VOLUME_BELONGS_TO_CONS_GROUP**

  The remote volume belongs to a consistency group.
- **TARGET_NOT_CONNECTED**

  There is currently no connection to the target system.
- **VOLUME_BAD_POOL**

  The volume belongs to a different storage pool.
- **VOLUME_BELONGS_TO_CG**

  The volume belongs to a consistency group.
- **VOLUME_DATA_MIGRATION_UNSYNCHRONIZED**

  Data Migration to this volume has not completed.
- **VOLUME_IS_SNAPSHOT**

  The operation is not permitted on snapshots.
- **CONS_GROUP_MIRROR_SCHEDULE_MISMATCH**

  All volumes in a mirrored consistency group must have the same mirroring schedule.
- **CONS_GROUP_MIRROR_TARGET_MISMATCH**

  All volumes in a mirrored consistency group must have the same mirroring target.
- **CONS_GROUP_MIRROR_ROLE_MISMATCH**

  All volumes in a mirrored consistency group must have the same mirroring role.
- **CONS_GROUP_MIRROR_ACTIVATION_MISMATCH**

  All volumes in a mirrored consistency group must have the same mirroring activation state.

- **CONS_GROUP_MIRROR_STANDBY_MISMATCH**

  All volumes in a mirrored consistency group must have the same mirroring standby state.
- **HA_HIGH_AVAILABILITY_DISABLED_IN_VOL**

  The consistency group's high availability is enabled, but the volume's high availability is disabled.
- **HA_HIGH_AVAILABILITY_ENABLED_IN_VOL**

  The consistency group's high availability is disabled but the volume's high availability is enabled.
- **CONS_GROUP_HA_ROLE_MISMATCH**

  All volumes in a HyperSwap consistency group must have the same mirroring role.
- **HA_LAST_SYNC_TIMES_DIFFER**

  All HyperSwap relations in a consistency group must have the same last sync time.
- **HA_POSSIBLE_CONS_GROUP_MEMBERSHIP_MISMATCH**

  The HyperSwap consistency group may contain different volumes on the primary and secondary machines.
- **CONS_GROUP_HA_ACTIVATION_MISMATCH**

  All volumes in a HyperSwap consistency group must have the same HyperSwap activation state.
- **CONS_GROUP_HA_TARGET_MISMATCH**

  All volumes in a mirrored consistency group must have the same HyperSwap target.
- **HA_RETRY_OPERATION**

  An operation is in progress on this HyperSwap relation.

  **Troubleshooting:** Try issuing the command again in a few seconds.
- **HA_IS_NOT_SYNCHRONIZED**

  The HyperSwap relation is not synchronized.
- **REMOTE_CONS_GROUP_MIRROR_SCHEDULE_MISMATCH**

  All volumes in a mirrored consistency group on the remote machine must have identical mirroring schedule.
- **CONS_GROUP_MIRROR_TYPE_MISMATCH**

  All volumes in a mirrored consistency group must be of the same mirroring type.
- **MIRROR_POSSIBLE_CONS_GROUP_MEMBERSHIP_MISMATCH**

  The mirrored consistency group contains different volumes on the primary and secondary machines. This problem occurs whenever the cg_add_vol command was issued, but the primary machine did not receive an acknowledgment from the secondary machine until the command timed out, or due to any other unexpected failure.
- **REMOTE_CONS_GROUP_CRASH_CONSISTENCY_MISMATCH**

  Crash consistency of the volume does not match the state of other volumes in the group on the remote machine.
- **CONS_GROUP_CRASH_CONSISTENCY_MISMATCH**

  Crash consistency of the volume does not match the state of other volumes in the group.
- **OPERATION_DENIED_OBJECT_MANAGED**

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

  An IBM Hyper-Scale Mobility relationship is defined for this volume.
- **VOLUME_HAS_MULTIPLE_RELATIONS**

The volume has multiple relations. The operation is not allowed or a target must be specified.

- **REMOTE_MIRROR_IS_STANDBY**

  The remote mirror is marked as Standby.
- **VOLUME_CANNOT_BE_ADDED_TO_CONS_GROUP_DUE_TO_CONS_GROUP_SIZE_LIMIT**

  Addition of the volume to the CG will cause it to exceed the maximal CG size.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support
- **MULTISITE_CONS_GROUP_CANNOT_CONTAIN_NON_MULTISITEED_VOLS**

  A Multi-site consistency group cannot include volumes that are not in a Multi-site relation.
- **MULTISITE_CONS_GROUP_ROLE_IS_NOT_MASTER**

  The local consistency group is not primary in the Multi-site relation.
- **MULTISITE_IS_NOT_OPERATIONAL**

  A Multi-site must be operational.
- **MULTISITE_IS_PARTIALY_DEFINED**

  The Multi-site relation is partially defined.
- **MULTISITE_CONS_GROUP_STANDBY_CONFIGURATION_MISMATCH**

  The standby state of the Multi-site consistency group does not match the standby configuration of a Multi-site volume.
- **MIRROR_OR_HYPERSWAP_CONS_GROUP_CANNOT_CONTAIN_MULTISITEED_VOLS**

  A mirroring or HyperSwap consistency group cannot include volumes that are in a Multi-site relation.
- **MULTISITE_ROLE_IS_STANDALONE_MASTER**

  Command failed because a Multi-site STANDALONE master cannot be activated.
- **MULTISITE_ROLE_IS_NOT_MASTER**

  The local peer is not the Master in the Multi-site relation.
- **MULTISITE_CONS_GROUP_LAST_SMASTER_SYNC_MISMATCH**

  The state of the Multi-site consistency group does not match the last SMaster sync of a Multi-site volume.
- **MULTISITE_CG_ADD_VOL_FAILED_TO_ROLLBACK_MANUAL_CLEANUP_REQUIRED**

  An error occured while adding a volume to the Multi-site consistency group, and the system failed to roll back. To recover, manually ensure that consistency group constituents are identical on all systems.

# Creating consistency groups

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

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

## Parameters

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

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

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

The storage pool of the consistency group must be specified.

The consistency group is initially empty, containing no volumes.

A consistency group always belongs to a specific storage pool. All the volumes in the consistency group belong to the same storage pool as the consistency group itself.

The consistency group can be mirrored as a whole (see Creating a mirroring definition).

**Example:**

```
cg_create pool=p_1 cg=DBgroup
```

**Output:**

```
Command completed successfully.
```

## Access control

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

## Return codes

- **CONS_GROUP_NAME_EXISTS**

  The consistency group name already exists.
- **MAX_CONS_GROUPS_REACHED**

  The maximum allowed number of consistency groups is already reached.
- **POOL_DOES_NOT_EXIST**

  The storage pool does not exist.
- **DOMAIN_MAX_CONS_GROUPS_REACHED**

  The domain exceeds the maximum allowed number of consistency groups.
- **OPERATION_DENIED_OBJECT_MANAGED**

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

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support

# Deleting a consistency group

Use the **cg_delete** command to delete a consistency group.

```
cg_delete cg=cgName
```

## Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **cg** | Object name | Name of the consistency group to be deleted. | Y |

This command fails if:

- The consistency group is not empty, that is, it still contains volumes.
- The consistency group is mirrored, even if it is empty.

All snapshot groups associated with the consistency group are disbanded. That means that the snapshots contained in these snapshot groups become independent snapshots.

**Example:**

```
cg_delete cg=DBvolumes
```

**Output:**

```
Command completed successfully
```

## Access control

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

## Return codes

- **CONS_GROUP_BAD_NAME**

  The consistency group name does not exist.
- **CONS_GROUP_NOT_EMPTY**

  This operation is only allowed on an empty consistency group.
- **CONS_GROUP_HAS_MIRROR**

  Mirroring is defined for this consistency group.
- **CONS_GROUP_BELONGS_TO_XCG**

  The consistency group belongs to another cross-system consistency group.
- **OPERATION_DENIED_OBJECT_MANAGED**

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

- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support

## Listing consistency groups

Use the **cg_list** command to list consistency groups.

```
cg_list [ cg=cgName ] [ managed=<yes|no|all> ] [ domain=DomainName ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **cg** | Object name | Name of a consistency group. | N | All |
| **managed** | Boolean | Determines whether to show unmanaged consistency groups (no), managed consistency groups (yes) or both (all). | N | no |
| **domain** | Object name | The domain name. | N | All Domains |

This command lists the specified details for all consistency groups. If a consistency group name is indicated, only this consistency group is listed.

The listed details include the following fields:

- **Name**
- **Mirrored** - Indicates whether the consistency group is mirrored
  - Available values: **Yes|No**
- **GP Based** - Indicates whether the consistency group is based on a grouped pool
  - Values - **Yes|No**
- **Mirror sync status** - Indicates the mirroring status
  - Available values: **RPO_OK|RPO_Lagging**
- **CG role** - Indicates the peer's role
  - Available values: **master|slave**

| Field ID | Field output | Default position |
|---|---|---|
| **name** | Name | 1 |
| **pool** | Pool Name | 2 |
| **mirrored** | Mirrored | N/A |
| **ha** | HA | N/A |
| **managed** | Managed | N/A |

**Example:**

```
cg_list cg=DBgroup
```

**Output:**

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

### Access control

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

## Removing a volume from a consistency group

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

```
cg_remove_vol vol=VolName
```

### Parameters

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

This command removes a volume from a consistency group.

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

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

Following the volume removal:

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

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

Requirements for a successful command completion:

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

If the command is issued on a mirrored contingency group master, and the master does not receive an acknowledgment from the slave because the command times out or due to an unexpected failure, a return code is returned: (**MIRROR_POSSIBLE_CONS_GROUP_MEMBERSHIP_MISMATCH**).

If applied on a volume that is part of an IBM Hyper-Scale Mobility relation:

- The command is applicable only to a destination volume.
- This destination volume has to be in Proxy state.

Otherwise, a completion code is returned (see below).

**Example:**

```
cg_remove_vol vol=DBLog
```

**Output:**

```
Command completed successfully
```

## Access control

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

## Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_REMOVE_VOLUME_FROM_CONS_GROUP**

  Are you sure you want to remove volume '*Volume*' from its consistency group?

## Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **VOLUME_NOT_IN_CONS_GROUP**

  The volume does not belong to a consistency group.
- **TARGET_NOT_CONNECTED**

  There is currently no connection to the target system.
- **VOLUME_IS_SNAPSHOT**

  The operation is not permitted on snapshots.
- **CONS_GROUP_IS_SLAVE**

  The consistency group's role in a mirroring relationship is secondary.
- **MIRROR_RETRY_OPERATION**

  There is an operation in progress on this mirror.

  **Troubleshooting:** Retry the command in a few seconds.
- **MIRROR_POSSIBLE_CONS_GROUP_MEMBERSHIP_MISMATCH**

  The mirrored consistency group contains different volumes on the primary and secondary machines. This problem occurs whenever the cg_add_vol command was issued, but the primary machine did not

receive an acknowledgment from the secondary machine until the command timed out, or due to any other unexpected failure.

- **HA_POSSIBLE_CONS_GROUP_MEMBERSHIP_MISMATCH**

  The HyperSwap consistency group may contain different volumes on the primary and secondary machines.

- **VOLUME_IS_NOT_CONSISTENT_SLAVE**

  The operation not allowed on an inconsistent secondary volume.

- **SNAPSHOT_HAS_ACTIVE_SYNC_JOB**

  The snapshot is currently the target of an active sync job.

  **Troubleshooting:** Please wait for the sync job to complete.

- **OPERATION_DENIED_OBJECT_MANAGED**

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

- **VOLUME_HAS_OLVM**

  An IBM Hyper-Scale Mobility relationship is defined for this volume.

- **REMOTE_MIRROR_IS_STANDBY**

  The remote mirror is marked as Standby.

- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support

- **MAX_SNAPSHOTS_PER_VOLUME_REACHED**

  The maximum allowed number of snapshots is already reached.

- **HA_RETRY_OPERATION**

  An operation is in progress on this HyperSwap relation.

  **Troubleshooting:** Try issuing the command again in a few seconds.

- **MIRROR_HAS_SYNC_JOB**

  The operation is not permitted on a mirror with active sync jobs.

## Renaming a consistency group

Use the **cg_rename** command to rename consistency groups.

```
cg_rename cg=cgName new_name=Name
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **cg** | Object name | The name of the consistency group to be renamed. | Y |
| **new_name** | Object name | The new name of the consistency group. | Y |

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

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

**Example:**

```
cg_rename cg=DBgroup new_name=DBvolumes
```

**Output:**

```
Command completed successfully
```

### Access control

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

### Return codes

- **CONS_GROUP_BAD_NAME**

  The consistency group name does not exist.
- **CONS_GROUP_NAME_EXISTS**

  The consistency group name already exists.
- **OPERATION_DENIED_OBJECT_MANAGED**

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

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support

## Creating a cross-system consistency group

Use the **xcg_create** command to create a cross-system consistency group (XCG) definition.

```
xcg_create xcg=XcgName
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **xcg** | Object name | The name of the new cross-system consistency group. | Y |

This command creates a cross-system consistency group (XCG) definition, with which consistency groups on different systems can be associated.

**Example:**

```
xcg_create xcg=DBbackup
```

**Output:**

```
Command completed successfully.
```

### Access control

| User Category | Permission | Condition |
|---|---|---|
| Storage administrator | Allowed | N/A |
| Storage integration administrator | Allowed | N/A |
| Application administrator | Conditionally Allowed | At least one of the volumes in the group is mapped to a host or cluster associated with the user. If a Snapshot Group overwrite is used, then the target Snapshot Group must be one created by a server administrator. |
| Security administrator | Disallowed | N/A |
| Read-only users | Disallowed | N/A |
| Technicians | Disallowed | N/A |

### Return codes

- **XCG_NAME_EXISTS**

  The cross-system consistency group name already exists.
- **MAX_XCGS_REACHED**

  The maximum allowed number of cross-system consistency groups is already reached.

## Associating an existing consistency group with a cross-system consistency group definition

Use the **xcg_add_cg** command to associate an existing consistency group to a cross-system consistency group definition.

```
xcg_add_cg xcg=XcgName cg=cgName
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **xcg** | Object name | Name of a cross-system consistency group. | Y |
| **cg** | Object name | Name of a consistency group. | Y |

**Example:**

```
xcg_add_cg xcg=DBbackup cg=CGbackup
```

**Output:**

```
Command completed successfully.
```

### Access control

| User Category | Permission | Condition |
|---|---|---|
| Storage administrator | Allowed | N/A |

| User Category | Permission | Condition |
|---|---|---|
| Storage integration administrator | Allowed | N/A |
| Application administrator | Conditionally Allowed | At least one of the volumes in the group is mapped to a host or cluster associated with the user. If a Snapshot Group overwrite is used, then the target Snapshot Group must be one created by a server administrator. |
| Security administrator | Disallowed | N/A |
| Read-only users | Disallowed | N/A |
| Technicians | Disallowed | N/A |

### Return codes

- **XCG_BAD_NAME**

  The cross-system consistency group name does not exist.
- **MAX_CONS_GROUPS_IN_XCG_REACHED**

  The cross-system consistency group contains the maximum allowed number of consistency groups.
- **CONS_GROUP_IS_SLAVE**

  The consistency group's role in a mirroring relationship is secondary.
- **CONS_GROUP_BAD_NAME**

  The consistency group name does not exist.
- **CONS_GROUP_ALREADY_IN_XCG**

  The consistency group already belongs to a cross-system consistency group.
- **CONS_GROUP_BELONGS_TO_XCG**

  The consistency group belongs to another cross-system consistency group.

## Removing an existing consistency group from a cross-system consistency group definition

Use the **xcg_remove_cg** command to remove an existing consistency group from a cross-system consistency group definition.

```
xcg_remove_cg xcg=XcgName cg=cgName
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **xcg** | Object name | Name of a Cross-system Consistency Group. | Y |
| **cg** | Object name | Name of a Consistency Group. | Y |

**Example:**

```
xcg_remove_cg xcg=DBbackup cg=CGBackup
```

**Output:**

```
Command completed successfully.
```

### Access control

| User Category | Permission | Condition |
|---|---|---|
| Storage administrator | Allowed | N/A |
| Storage integration administrator | Allowed | N/A |
| Application administrator | Conditionally Allowed | At least one of the volumes in the group is mapped to a host or cluster associated with the user. If a Snapshot Group overwrite is used, then the target Snapshot Group must be one created by a server administrator. |
| Security administrator | Disallowed | N/A |
| Read-only users | Disallowed | N/A |
| Technicians | Disallowed | N/A |

### Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_REMOVE_CONS_GROUP_FROM_XCG**

  Are you sure you want to remove consistency group '*CG*' from its cross-system consistency group?

### Return codes

- **XCG_BAD_NAME**

  The cross-system consistency group name does not exist.
- **CONS_GROUP_BAD_NAME**

  The consistency group name does not exist.
- **XCG_IS_EMPTY**

  The consistency group is empty.
- **CONS_GROUP_NOT_IN_XCG**

  The consistency group does not belong to a cross-system consistency group.

## Adding a remote system name to a cross-system consistency group definition

Use the **xcg_add_remote_system** command to add a remote system name to a cross-system consistency group definition.

```
xcg_add_remote_system xcg=XcgName remote_system=RemoteSystem
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **xcg** | Object name | Name of a cross-system consistency group. | Y |
| **remote_system** | String | Name of a remote system. | Y |

**Example:**

```
xcg_add_remote_system xcg=DBbackup remote_system=CGbackup
```

**Output:**

```
Command completed successfully.
```

### Access control

| User Category | Permission | Condition |
|---|---|---|
| Storage administrator | Allowed | N/A |
| Storage integration administrator | Allowed | N/A |
| Application administrator | Conditionally Allowed | At least one of the volumes in the group is mapped to a host or cluster associated with the user. If a Snapshot Group overwrite is used, then the target Snapshot Group must be one created by a server administrator. |
| Security administrator | Disallowed | N/A |
| Read-only users | Disallowed | N/A |
| Technicians | Disallowed | N/A |

### Return codes

- **XCG_BAD_NAME**

  The cross-system consistency group name does not exist.
- **MAX_REMOTE_SYSTEMS_IN_XCG_REACHED**

  The cross-system consistency group contains the maximum number of remote systems.
- **REMOTE_SYSTEM_ALREADY_ADDED**

  The remote system belongs to a cross-system consistency group.

## Removing a remote system name from a cross-system consistency group definition

Use the **xcg_remove_remote_system** command to remove a remote system name from a cross-system consistency group definition.

```
xcg_remove_remote_system xcg=XcgName remote_system=RemoteSystem
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **xcg** | Object name | Name of a Cross-system Consistency Group. | Y |
| **remote_system** | String | Name of a remote system. | Y |

**Example:**

```
xcg_remove_remote_system xcg=DBbackup remote_system=CGbackup
```

**Output:**

```
Command completed successfully.
```

### Access control

| User Category | Permission | Condition |
|---|---|---|
| Storage administrator | Allowed | N/A |
| Storage integration administrator | Allowed | N/A |
| Application administrator | Conditionally Allowed | At least one of the volumes in the group is mapped to a host or cluster associated with the user. If a Snapshot Group overwrite is used, then the target Snapshot Group must be one created by a server administrator. |
| Security administrator | Disallowed | N/A |
| Read-only users | Disallowed | N/A |
| Technicians | Disallowed | N/A |

### Return codes

- **XCG_BAD_NAME**

  The cross-system consistency group name does not exist.
- **REMOTE_SYSTEM_NOT_IN_XCG**

  The remote system does not belong to a cross-system consistency group.

## Listing cross-system consistency group definitions

Use the **xcg_get_local_cgs** command to list cross-system consistency group definitions together with the contained consistency groups.

```
xcg_get_local_cgs [ xcg=XcgName ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **xcg** | Object name | Name of a cross-system consistency group. | N | All Cross-system Consistency Groups. |

**Example:**

```
xcg_get_local_cgs
```

**Output:**

```
Command completed successfully.
```

| Field ID | Field output | Default position |
|---|---|---|
| **name** | Name | 1 |
| **xcg** | XCG Name | 2 |

### Access control

| User Category | Permission | Condition |
|---|---|---|
| Storage administrator | Allowed | N/A |
| Storage integration administrator | Allowed | N/A |

| User Category | Permission | Condition |
|---|---|---|
| Application administrator | Conditionally Allowed | At least one of the volumes in the group is mapped to a host or cluster associated with the user. If a Snapshot Group overwrite is used, then the target Snapshot Group must be one created by a server administrator. |
| Security administrator | Disallowed | N/A |
| Read-only users | Allowed | N/A |
| Technicians | Disallowed | N/A |

### Return codes

- **XCG_BAD_NAME**

  The cross-system consistency group name does not exist.

## Retrieving remote systems in a specified cross-system consistency group

Use the **xcg_get_remote_systems** command to retrieve the names of remote systems that are a part of the specified cross-system consistency group.

```
xcg_get_remote_systems xcg=XcgName
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **xcg** | Object name | Name of a Cross-system Consistency Group. | Y |

**Example:**

```
xcg_get_remote_systems xcg=blabla
```

**Output:**

```
Command completed successfully.
```

| Field ID | Field output | Default position |
|---|---|---|
| **name** | Name | 1 |
| **xcg** | XCG Name | 2 |

### Access control

| User Category | Permission | Condition |
|---|---|---|
| Storage administrator | Allowed | N/A |
| Storage integration administrator | Allowed | N/A |
| Application administrator | Conditionally Allowed | At least one of the volumes in the group is mapped to a host or cluster associated with the user. If a Snapshot Group overwrite is used, then the target Snapshot Group must be one created by a server administrator. |

| User Category | Permission | Condition |
|---|---|---|
| Security administrator | Disallowed | N/A |
| Read-only users | Allowed | N/A |
| Technicians | Disallowed | N/A |

### Return codes

- **XCG_BAD_NAME**

  The cross-system consistency group name does not exist.

## Deleting a cross-system consistency group

Use the **xcg_delete** command to delete a cross-system consistency group (XCG) definition.

```
xcg_delete xcg=XcgName
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **xcg** | Object name | Name of a cross-system consistency group. | Y |

**Example:**

```
xcg_delete xcg=DBbackup
```

**Output:**

```
Command completed successfully.
```

### Access control

| User Category | Permission | Condition |
|---|---|---|
| Storage administrator | Allowed | N/A |
| Storage integration administrator | Allowed | N/A |
| Application administrator | Conditionally Allowed | At least one of the volumes in the group is mapped to a host or cluster associated with the user. If a Snapshot Group overwrite is used, then the target Snapshot Group must be one created by a server administrator. |
| Security administrator | Disallowed | N/A |
| Read-only users | Disallowed | N/A |
| Technicians | Disallowed | N/A |

### Return codes

- **XCG_BAD_NAME**

  The cross-system consistency group name does not exist.
- **XCG_NOT_EMPTY**

  The consistency group is not empty.

# Listing cross-system consistency group definitions

Use the **xcg_list** command to list cross-system consistency group definitions together with the contained consistency groups.

```
xcg_list [ xcg=XcgName ]
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **xcg** | Object name | Name of a Cross-system Consistency Group. | N | All Cross-system Consistency Groups. |

| Field ID | Field output | Default position |
|---|---|---|
| **name** | Name | 1 |
| **num_of_cgs** | Num Of CGs | 2 |
| **num_of_remote_systems** | Num Of Remote Systems | 3 |

**Example:**

```
xcg_list
```

**Output:**

```
Command completed successfully.
```

## Access control

| User Category | Permission | Condition |
|---|---|---|
| Storage administrator | Allowed | N/A |
| 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.
- 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.

© Copyright IBM Corp. 2016, 2022

- All I/Os to the consistency group's volumes that were completed prior to this point of time are recorded in the snapshot's image.
- Neither I/O that was completed after this point of time is recorded in the snapshot's image.

In addition to their regular attributes, all the snapshots in the snapshot group are also associated with the consistency group.

The name of the snapshot group is either automatically generated or provided in the command line.

The delete priority of the snapshots in the snapshot group can also be provided (see Creating a snapshot). The delete priority controls which snapshots or snapshot groups are deleted first when the system runs out of space for snapshots.

The `overwrite` option causes the current content of the consistency group to be copied into one of its existing snapshot groups (indicated as parameter's argument). The snapshots of the overwritten snapshot group keep the same SCSI device WWN and same mapping, so hosts maintain a continuous mapping of the snapshots, and a rescan or similar operation is not needed. The overwritten snapshot group must be an existing snapshot group of the respective consistency group.

This command fails if no snapshot space is defined for the storage pool containing the consistency group.

This command fails if one or more of the volumes in the consistency group are slaves in the synchronous mirroring, and the synchronous mirroring is currently inconsistent due to either a re-synchronization or an initialization process.

Mirroring limitations:

- This command fails if the volume is a slave of an asynchronous mirroring coupling (either synchronous or asynchronous).
- This command fails if the volume is a slave of an inconsistent synchronous coupling.

**Example:**

```
cg_snapshots_create cg=DBgroup snap_group=DBbackupdaily
```

**Output:**

```
Command completed successfully.
```

## Access control

| User Category | Permission | Condition |
|---|---|---|
| Storage administrator | Allowed | N/A |
| Storage integration administrator | Allowed | N/A |
| Application administrator | Conditionally Allowed | At least one of the volumes in the group is mapped to a host or cluster associated with the user. If a Snapshot Group overwrite is used, then the target Snapshot Group must be one created by a server administrator. |
| Security administrator | Disallowed | N/A |
| Read-only users | Disallowed | N/A |
| Technicians | Disallowed | N/A |

## Return codes

- **MAX_VOLUMES_REACHED**

  The maximum allowed number of volumes is already reached.

- **DOMAIN_MAX_VOLUMES_REACHED**

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

  The consistency group name does not exist.
- **SNAPSHOT_GROUP_BAD_NAME**

  The snapshot group name does not exist.
- **SNAPSHOT_GROUP_BAD_PREFIX**

  The snapshot group name has a reserved prefix.
- **SNAPSHOT_GROUP_NAME_EXISTS**

  The snapshot group name already exists.
- **CONS_GROUP_EMPTY**

  The operation is not allowed on an empty consistency group.
- **CONS_GROUP_MISMATCH**

  The snapshot group does not match the consistency group volumes.
- **OVERWRITE_SNAPSHOT_GROUP_DOES_NOT_BELONG_TO_GIVEN_GROUP**

  The snapshot group belongs to another consistency group.
- **POOL_SNAPSHOT_LIMIT_REACHED**

  There is not enough space to create a snapshot.
- **VOLUME_IS_NOT_CONSISTENT_SLAVE**

  The operation not allowed on an inconsistent secondary volume.
- **SNAPSHOT_GROUP_IS_INTERNAL**

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

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

  The snapshot is currently the target of an active sync job.

  **Troubleshooting:** Please wait for the sync job to complete.
- **OPERATION_DENIED_OBJECT_MANAGED**

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

  The token does not match the consistency group.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support
- **MAX_SNAPSHOTS_PER_VOLUME_REACHED**

  The maximum allowed number of snapshots is already reached.

# Changing a snapshot group deletion priority

Use the **snap_group_change_priority** command to change the deletion priority of a snapshot group.

```
snap_group_change_priority snap_group=SnapshotGroupName  delete_priority=del_value
```

## Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **snap_group** | Object name | Name of the snapshot group whose delete_priority is to be changed. | Y |
| **delete_priority** | Integer | Priority according to which this snapshot group is deleted. | Y |

This command changes the priority of the deletion of an existing snapshot group. Similarly to snapshots, the system determines which of the snapshot groups is deleted first when it runs out of snapshot storage, in accordance with the redirect-on-write mechanism. When the system runs out of space, it deletes the snapshot or snapshot group with the highest deletion priority, and among them the unmapped snapshots or snapshot groups, and the snapshot or snapshot group which was created first.

See Changing a snapshot deletion priority for more details about the valid deletion priority values and their meaning.

**Example:**

```
snap_group_change_priority snap_group=DBbackup delete_priority=4
```

**Output:**

```
Command completed successfully.
```

## Access control

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

## Return codes

- **SNAPSHOT_GROUP_BAD_NAME**

  The snapshot group name does not exist.
- **SNAPSHOT_ILLEGAL_PRIORITY**

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

Internal snapshots cannot be mapped, modified, or deleted.

- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support

# Deleting a snapshot group

Use the **snap_group_delete** command to delete a snapshot group and all its snapshots.

```
snap_group_delete snap_group=SnapshotGroupName
```

## Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **snap_group** | Object name | Name of the snapshot group to be deleted. | Y |

This command deletes the snapshot group, as well as all of the snapshots that are contained in the snapshot group. Refer to the documentation on Deleting a snapshot for more information about deleting snapshots.

If one of the members of the snapshot group is mapped to a host, then the entire snapshot group cannot be deleted.

The command is inapplicable for a snapshot group that is still associated with a mirrored consistency group.

**Example:**

```
snap_group_delete snap_group=DBBackupweekly
```

**Output:**

```
Command completed successfully.
```

## Access control

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

## Return codes

- **OPERATION_DENIED_OBJECT_MANAGED**

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

- **SNAPSHOT_GROUP_BAD_NAME**

  The snapshot group name does not exist.
- **SNAPSHOT_IS_MAPPED**

  A snapshot that is mapped to a host cannot be deleted.
- **VOLUME_IS_BOUND**

  The volume is bound to an ALU.

  **Troubleshooting:** Unbind the volume from the ALU.
- **SNAPSHOT_HAS_ACTIVE_SYNC_JOB**

  The snapshot is currently the target of an active sync job.

  **Troubleshooting:** Please wait for the sync job to complete.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support

# Disbanding a snapshot group

Use the **snap_group_disband** command to disband a snapshot group into independent snapshots.

```
snap_group_disband snap_group=SnapshotGroupName
```

## Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **snap_group** | Object name | Snapshot group to be disbanded. | Y |

This command disbands the snapshot group into independent snapshots. After executing this command, the snapshots can be individually deleted, restored, unlocked, duplicated, and so on. The snapshot group does not exist anymore after this command. The snapshots retain the same names (**snap_group_name.volumename**).

The command is inapplicable for a snapshot group of a mirrored consistency group.

**Example:**

```
snap_group_disband snap_group=DBbackup_copy
```

**Output:**

```
Command completed successfully.
```

## Access control

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

| User Category | Permission | Condition |
|---|---|---|
| Security administrator | Disallowed | N/A |
| Read-only users | Disallowed | N/A |
| Technicians | Disallowed | N/A |

### Return codes

- **OPERATION_DENIED_OBJECT_MANAGED**

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

  The snapshot group name does not exist.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support

## Duplicating a snapshot group

Use the **snap_group_duplicate** command to duplicate an existing snapshot group.

```
snap_group_duplicate snap_group=SnapshotGroupName [ new_snap_group=NewName ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **snap_group** | Object name | Name of the snapshot group to be duplicated. | Y | N/A |
| **new_snap_group** | Object name | Name of the newly generated snapshot group. | N | Autogenerated name. |

This command duplicates the specified snapshot group. This is functionally equivalent to duplicating all the snapshots in the snapshot group using Duplicating a snapshot and creating a new snapshot group that contains all the generated snapshots.

The name of the new snapshot group is either specified as a parameter or generated automatically.

Refer to Duplicating a snapshot for more details about the snapshot duplication operation.

Deletion priority:

- The deletion priority of the duplicated snapshots is 0.

**Example:**

```
snap_group_duplicate snap_group=DBbackup new_snap_group=DBbackup_copy
```

**Output:**

```
Command completed successfully
```

### Access control

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

### Return codes

- **SNAPSHOT_GROUP_BAD_NAME**

  The snapshot group name does not exist.
- **MAX_VOLUMES_REACHED**

  The maximum allowed number of volumes is already reached.
- **DOMAIN_MAX_VOLUMES_REACHED**

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

  The snapshot group name already exists.
- **OPERATION_DENIED_OBJECT_MANAGED**

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

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support
- **MAX_SNAPSHOTS_PER_VOLUME_REACHED**

  The maximum allowed number of snapshots is already reached.

## Formatting a snapshot group

Use the **snap_group_format** command to format a snapshot group.

```
snap_group_format  snap_group=SnapshotGroupName
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **snap_group** | Object name | The snapshot group to be formatted. | Y |

This command deletes the content of a snapshot group while maintaining its snapshots mapping to the host. The format operation results with:

- The snapshots of the formatted snapshot group are read-only

- The format operation has no impact on performance
- The snapshots of the formatted snapshot group do not consume space
- Reading from the snapshots of the formatted snapshot group always returns zeroes
- The snapshots can be overridden
- The snapshots can be deleted
- The snapshots deletion priority can be changed

**Example:**

```
snap_group_format snap_group
```

**Output:**

```
Command completed successfully.
```

## Access control

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

## Return codes

- **SNAPSHOT_HAS_ACTIVE_SYNC_JOB**

  The snapshot is currently the target of an active sync job.

  **Troubleshooting:** Please wait for the sync job to complete.
- **SNAPSHOT_GROUP_IS_FORMATTED**

  The snapshot group is formatted.
- **ELCS_GROUP_CANNOT_BE_FORMATTED**

  The snapshot group is an ELCS (external last consistent snapshot), and cannot be formatted.
- **MAX_VOLUMES_REACHED**

  The maximum allowed number of volumes is already reached.
- **SNAPSHOT_GROUP_IS_INTERNAL**

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

  The operation is permitted on snapshots only.
- **SNAPSHOT_GROUP_BAD_NAME**

  The snapshot group name does not exist.
- **OPERATION_DENIED_OBJECT_MANAGED**

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

- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support

## Listing snapshot groups

Use the **snap_group_list** command to list all snapshot groups or a specific one.

```
snap_group_list [ snap_group=SnapshotGroupName | cg=cgName ] [ managed=<yes|no|all> ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **snap_group** | Object name | Name of a specific snapshot group to be listed. | N | All snapshot groups. |
| **cg** | Object name | List all the snapshot groups of this Consistency Group. | N | All snapshot groups. |
| **managed** | Boolean | Defines whether to show unmanaged snap groups (no), managed (yes) or both (all). | N | no. |

This command lists snapshot groups. When a snapshot group name is specified, then only that specific snapshot group is listed. When a consistency group name is specified, then the snapshot groups of this consistency group are listed.

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

- **snap_group_format**

| Field ID | Field output | Default position |
|---|---|---|
| **name** | Name | 1 |
| **cg** | CG | 2 |
| **snapshot_time** | Snapshot Time | 3 |
| **locked** | Locked | N/A |
| **modified** | Modified | N/A |
| **delete_priority** | Deletion Priority | 4 |
| **snap_group_format** | Snapshot Group Format | N/A |
| **snap_group_descriptor** | Snapshot Group Descriptor | N/A |
| **managed** | Managed | N/A |

**Example:**

```
snap_group_list cg=DBvolumes
```

**Output:**

```
Name           CG         Snapshot Time       Deletion Priority
DBbackup       DBvolumes  2007-01-03 17:46:29 1
DBbackupdaily  DBvolumes  2007-01-03 17:49:36 1
```

### Access control

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

## Locking a snapshot group

Use the **snap_group_lock** command to lock a snapshot group by locking all its snapshots.

```
snap_group_lock snap_group=SnapshotGroupName
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **snap_group** | Object name | Name of the snapshot group to be locked. | Y |

This command is functionally equivalent to locking all snapshots individually (through executing Locking a volume on each snapshot). Refer to the documentation of Locking a volume for a description of locking behavior.

**Example:**

```
snap_group_lock snap_group=DBbackup
```

**Output:**

```
Command completed successfully.
```

### Access control

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

### Return codes

- **SNAPSHOT_GROUP_BAD_NAME**

  The snapshot group name does not exist.
- **SNAPSHOT_GROUP_IS_INTERNAL**

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

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support

## Renaming a snapshot group

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

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

### Parameters

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

**Example:**

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

**Output:**

```
Command completed successfully
```

### Access control

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

### Return codes

- **SNAPSHOT_GROUP_BAD_NAME**

  The snapshot group name does not exist.
- **SNAPSHOT_GROUP_NAME_EXISTS**

The snapshot group name already exists.

- **OPERATION_DENIED_OBJECT_MANAGED**

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

- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support

## Restoring a consistency group from a snapshot group

Use the **snap_group_restore** command to restore the master volumes of a consistency group, or of a snapshot group from one of its associated snapshot groups.

```
snap_group_restore snap_group=SnapshotGroupName [ target_snap_group=SnapshotGroupName ]
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **snap_group** | Object name | Name of the snapshot group from which to restore its master volumes. | Y |
| **target_snap_group** | Object name | Snapshot group to be restored. | N |

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

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

Requirements for a successful command completion:

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

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

**Example:**

```
snap_group_restore snap_group=DBbackup_copy
```

**Output:**

```
Command completed successfully.
```

### Access control

| User Category | Permission | Condition |
|---|---|---|
| Storage administrator | Allowed | N/A |
| Storage integration administrator | Allowed | N/A |

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

## Return codes

- **VOLUME_HAS_DATA_MIGRATION**

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

  The snapshot group name does not exist.
- **CONS_GROUP_MISMATCH**

  The snapshot group does not match the consistency group volumes.
- **VOLUME_TOO_BIG**

  No space to allocate to the volume.
- **VOLUME_HAS_MIRROR**

  A mirror is defined for this volume.
- **VOLUME_HAS_HA**

  This operation is forbidden on a volume with a HyperSwap relation.
- **CONS_GROUP_HAS_MIRROR**

  Mirroring is defined for this consistency group.
- **VOLUME_LOCKED**

  The volume is locked.
- **TARGET_SNAPSHOT_GROUP_BAD_NAME**

  The target snapshot group name does not exist.
- **SNAPSHOT_GROUP_MISMATCH**

  The snapshot group does not match the target snapshot group.
- **TARGET_SNAPSHOT_GROUP_SAME_AS_SOURCE**

  The target snapshot group is identical with the snapshot group.
- **OPERATION_DENIED_OBJECT_MANAGED**

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

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support

# Unlocking a snapshot group

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

```
snap_group_unlock snap_group=SnapshotGroupName
```

## Parameters

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

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

**Example:**

```
snap_group_unlock snap_group=DBbackup
```

**Output:**

```
Command completed successfully
```

## Access control

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

## Return codes

- **SNAPSHOT_GROUP_BAD_NAME**

  The snapshot group name does not exist.
- **SNAPSHOT_GROUP_IS_INTERNAL**

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

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

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support

# Setting a snapshot group descriptor

Use the **snap_group_set_descriptor** command to set a snapshot group descriptor.

```
snap_group_set_descriptor snap_group=SnapshotGroupName descriptor=Descriptor
```

## Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **snap_group** | Object name | Name of the snapshot group. | Y |
| **descriptor** | String | A snap group descriptor to be used by external software. | Y |

Provides external software with the ability to mark the snapshot as part of a consistency group for various usage scenarios. The command replaces an existing descriptor with a newly specified one.

**Example:**

```
snap_group_set_descriptor snap_group=DBbackup descriptor=blabla
```

**Output:**

```
Command completed successfully
```

## Access control

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

## Return codes

- **SNAPSHOT_GROUP_BAD_NAME**

  The snapshot group name does not exist.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support

# Returning a snapshot group's descriptor

Use the **snap_group_get_descriptor** command to return a snapshot group's descriptor.

```
snap_group_get_descriptor snap_group=SnapshotGroupName
```

## Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **snap_group** | Object name | Name of the snapshot group. | Y |

The command provides an external software with the ability to obtain the descriptor attribute value for a snapshot group.

**Example:**

```
snap_group_get_descriptor snap_group=DBbackup
```

**Output:**

```
Command completed successfully
```

## Access control

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

## Return codes

- **SNAPSHOT_GROUP_BAD_NAME**

  The snapshot group name does not exist.

# 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. If the consistency group is mirrored, it can only be moved to a pool that is not thin-provisioned.

**Example:**

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

**Output:**

```
Command completed successfully.
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Allowed |
| Application administrator | Disallowed |
| Security administrator | Disallowed |

© Copyright IBM Corp. 2016, 2022

| User Category | Permission |
|---|---|
| Read-only users | Disallowed |
| Technicians | Disallowed |

## Return codes

- **CONS_GROUP_BAD_NAME**

  The consistency group name does not exist.
- **POOL_DOES_NOT_EXIST**

  The storage pool does not exist.
- **NOT_ENOUGH_SPACE**

  No space to allocate for the volume's current usage.
- **NOT_ENOUGH_SNAPSHOT_SPACE**

  Snapshot usage will exceed the snapshot limit.
- **DOMAIN_MAX_VOLUMES_REACHED**

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

  The maximum allowed number of volumes is already reached.
- **DOMAIN_MAX_CONS_GROUPS_REACHED**

  The domain exceeds the maximum allowed number of consistency groups.
- **MAX_CONS_GROUPS_REACHED**

  The maximum allowed number of consistency groups is already reached.
- **DOMAIN_MAX_MIRRORS_REACHED**

  The domain exceeds the maximum allowed number of mirrors.
- **MAX_MIRRORS_REACHED**

  The maximum number of mirrors is already reached.
- **DOMAIN_USED_TARGET_NOT_IN_DESTINATION**

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

  A schedule used by a mirror in the pool is not associated with the target domain.
- **MAPPED_HOSTS_NOT_IN_DESTINATION**

  A host that is mapped to a volume in the pool is not associated with the target domain.
- **MAPPED_CLUSTERS_NOT_IN_DESTINATION**

  A cluster that is mapped to a volume in the pool is not associated with the target domain.
- **OPERATION_DENIED_OBJECT_MANAGED**

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

  A destination pool must be defined.
- **MAX_DMS_REACHED**

  The maximum number of remote volumes (mirror/migration) is already reached.

  **Troubleshooting:** Delete unnecessary data migration objects.

- **DOMAIN_MAX_DMS_REACHED**

  The domain exceeds the maximum allowed number of data migrations.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support
- **NO_SPACE**

  The system does not have enough free space for the requested storage pool size.
- **VOLUME_TOO_BIG**

  No space to allocate to the volume.

# Changing the pool limitation, performance class, or threshold parameters

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

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

## Parameters

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

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

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

**Example:**

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

This command changes the Performance Class of the pool.

**Example:**

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

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

**Example:**

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

**Output:**

```
Command completed successfully.
```

## Access control

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

## Return codes

- **POOL_DOES_NOT_EXIST**

  The storage pool does not exist.
- **PERF_CLASS_BAD_NAME**

  The performance class does not exist.
- **PERF_CLASS_ASSOCIATED_WITH_HOSTS**

  Performance class *Performance Class* is already being used by a host.
- **PERF_CLASS_ASSOCIATED_WITH_VOLUMES**

  Performance class *Performance Class* is already being used by a volume.
- **POOL_ALREADY_IN_PERF_CLASS**

  Pool *pool name* is already in performance class *Performance Class*.
- **OPERATION_DENIED_OBJECT_MANAGED**

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

  '*String*' is not a recognized return code.

  **Troubleshooting:** Consult the manual for the list of valid return codes.
- **EVENT_DOES_NOT_HAVE_THRESHOLDS**

  The event does not have thresholds.
- **EVENT_THRESHOLD_IS_ILLEGAL**

  An illegal value for the event threshold.

**Troubleshooting:** Event threshold values must be monotonic.

# Changing pool settings for snapshots

Use the **pool_config_snapshots** command to change storage pool snapshot settings.

```
pool_config_snapshots pool=PoolName [ protected_snapshot_priority=<0|1|2|3|4> ]
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **pool** | Object name | The name of a storage pool. | Y | N/A |
| **protected_snapshot_priority** | Integer | Specifies the snapshot delete priority from 0 to 4 (see full explanation below). | N | unchanged |

This command changes the storage pool snapshot limitation policy.

The *create_last_consistent_snapshot* attribute (used for systems which have no space):

- If the value of the attribute is No, no last consistent snapshot is generated.
- If the value is changed while synchronizing, the existing snapshot is not deleted.

The **protected_snapshot_priority** parameter:

- Snapshots with a lower delete priority (that is, a higher number) than the specified value might be deleted by the system automatically, in order to free space, before pausing the mirroring, thus protecting snapshots with a priority equal or higher than the value.
- If, for example, the value is set to 3:
  - The system will deactivate mirroring if not enough space can be freed even after the deletion of snapshots with deletion priority of 4.
  - Snapshots with priority level 1, 2 and 3 will not be deleted.
- If the value is set to 4, the system will deactivate mirroring before deleting any of the snapshots.
- If the value is set to 0, the system can delete any snapshot regardless of deletion priority.

**Example:**

```
pool_config_snapshots pool=DBPool
```

**Output:**

```
Command completed successfully.
```

## Access control

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

### Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_CHANGE_THE_PROTECTED_LEVEL_OF_SNAPSHOTS**

  Are you sure you want to change the protection level of a snapshot in storage pool *Pool*? Note that in case of pool space depletion the system will delete protected snapshots only after deleting unprotected snapshots and internal asynchronous mirror snapshots.
- **ARE_YOU_SURE_YOU_WANT_TO_INCREASE_THE_PROTECTED_LEVEL_OF_EXISTING_SNAPSHOTS**

  Are you sure you want to increase the protection level of a snapshot in storage pool *Pool*? Note that the pool contains unprotected snapshots that will become protected after issuing this command. In case of pool space depletion the system will delete protected snapshots only after deleting unprotected snapshots and internal asynchronous mirror snapshots.
- **ARE_YOU_SURE_YOU_WANT_TO_DECREASE_THE_PROTECTED_LEVEL_OF_EXISTING_SNAPSHOTS**

  Are you sure you want to decrease the protection level of a snapshot in Storage Pool *Pool*? Note that the pool contains protected snapshots that will become unprotected after issuing this command. In case of pool space depletion the system will delete internal asynchronous mirror snapshots only after deleting unprotected snapshots.

### Return codes

- **POOL_DOES_NOT_EXIST**

  The storage pool does not exist.
- **SNAPSHOT_ILLEGAL_PRIORITY**

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

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

## Creating storage pools

Use the **pool_create** command to create a storage pool.

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

### Parameters

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

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

**Pool size limits**

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

The upper limit of the **size** parameter is not restricted, provided that the total effective capacity of all pools does not exceed the storage system's maximum effective capacity. However, to protect the system against overallocation of its resources, the maximum allowed total effective capacity of the volumes in a consistency group is limited to the effective capacity of the system.

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

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

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

**Example:**

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

**Output:**

```
Command completed successfully.
```

## Access control

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

## Warnings

- **POOL_SIZE_SMALL**

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

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

## Return codes

- **POOL_NAME_EXISTS**

  The storage pool name already is assigned to another storage pool.
- **PERF_CLASS_BAD_NAME**

  The performance class does not exist.
- **PERF_CLASS_ASSOCIATED_WITH_HOSTS**

  Performance class *Performance Class* is already being used by a host.
- **PERF_CLASS_ASSOCIATED_WITH_VOLUMES**

Performance class *Performance Class* is already being used by a volume.

- **MAX_POOLS_REACHED**

  The maximum allowed number of storage pools is already reached.
- **NO_SPACE**

  The system does not have enough free space for the requested storage pool size.
- **SNAPSHOT_SIZE_BIGGER_THAN_POOL_SIZE**

  The snapshot size must be equal to or smaller than the pool size.
- **REACHED_POOL_MAX_SIZE**

  Maximum pool size usage is already reached.
- **DOMAIN_DOESNT_EXIST**

  The domain does not exist.
- **USER_ASSOCIATED_TO_MORE_THAN_ONE_DOMAIN**

  The current user is attached to more than one domain, and it is not clear in which domain the pool is to be created.

  **Troubleshooting:** Re-run the command by specifying a domain.
- **NO_FREE_CAPACITY_IN_DOMAIN**

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

  The maximum allowed number of domain pools is already reached.

# Deleting a storage pool

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

```
pool_delete pool=PoolName
```

## Parameters

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

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

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

**Example:**

```
pool_delete pool=ERPPool
```

**Output:**

```
Command completed successfully
```

## Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Allowed |

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

### Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_DELETE_POOL**

  Are you sure you want to delete storage pool *Pool*?

### Return codes

- **POOL_DOES_NOT_EXIST**

  The storage pool does not exist.
- **POOL_HAS_CG**

  The storage pool comprises consistency groups.
- **POOL_IN_USE**

  The storage pool comprises allocated volumes.
- **OPERATION_DENIED_OBJECT_MANAGED**

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

## Listing storage pools

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

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

### Parameters

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

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **new_name** | Object name | The new name of the storage pool. | Y |

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

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

**Example:**

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

**Output:**

```
Command completed successfully.
```

### Access control

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

### Return codes

- **POOL_DOES_NOT_EXIST**

  The storage pool does not exist.
- **POOL_NAME_EXISTS**

  The storage pool name already is assigned to another storage pool.
- **OPERATION_DENIED_OBJECT_MANAGED**

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

## Resizing a storage pool

Use the **pool_resize** command to resize a storage pool.

```
pool_resize pool=PoolName [ size=GB ] [ snapshot_size=GB ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **pool** | Object name | The name of the storage pool to be resized. | Y | N/A |
| **size** | Positive integer | The new size of the storage pool (in gigabytes) | N | N/A |

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **snapshot_size** | Integer | The new limit on snapshot capacity usage of the storage pool. | N | Leave unchanged. |

The command can either increase or decrease the storage pool size.

The new size of the storage pool is specified as an integer multiple of $10^9$ bytes, but the actual size of the created storage pool is rounded up to the nearest integer multiple of $16x2^{30}$ bytes.

Capacity accounting is performed relative to the free space.

You can either specify both hard size and soft size or one size only, which assumes that the hard size and the soft size are identical.

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

Pool size limits

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

The upper of the size parameter is set to the effective system capacity.

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

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

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

This command fails if the current storage pool size (hard or soft) cannot be decreased, or if the free space (hard or soft) cannot be decreased.

Resizing a pool that stores asynchronously mirrored volumes or consistency groups:

- The pool's hard and soft sizes must be identical.

**Example:**

```
pool_resize pool=DBPool size=1300
```

**Output:**

```
Command completed successfully.
```

## Access control

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

### Warnings

- **POOL_SIZE_SMALL**

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

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

### Return codes

- **POOL_DOES_NOT_EXIST**

  The storage pool does not exist.
- **POOL_SIZE_TOO_SMALL**

  Storage pool usage exceeds the requested size.
- **REACHED_POOL_MAX_SIZE**

  Maximum pool size usage is already reached.
- **NO_SPACE**

  The system does not have enough free space for the requested storage pool size.
- **POOL_SNAPSHOT_SIZE_TOO_SMALL**

  Storage pool snapshot usage exceeds the requested snapshot size.
- **SNAPSHOT_SIZE_BIGGER_THAN_POOL_SIZE**

  The snapshot size must be equal to or smaller than the pool size.
- **OPERATION_DENIED_OBJECT_MANAGED**

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

  There is not enough free space in the domain.

## Moving a volume between storage pools

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

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

### Parameters

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

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

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

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

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

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

**Example:**

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

**Output:**

```
Command completed successfully.
```

## Access control

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

## Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **POOL_DOES_NOT_EXIST**

  The storage pool does not exist.
- **NOT_ENOUGH_SPACE**

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

  The system does not have enough free space for the requested storage pool size.
- **VOLUME_TOO_BIG**

  No space to allocate to the volume.
- **NO_FREE_CAPACITY_IN_DOMAIN**

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

  The operation is not permitted on snapshots.
- **VOLUME_HAS_OLVM**

An IBM Hyper-Scale Mobility relationship is defined for this volume.

- **VOLUME_BELONGS_TO_CG**

  The volume belongs to a consistency group.

- **NOT_ENOUGH_SNAPSHOT_SPACE**

  Snapshot usage will exceed the snapshot limit.

- **OPERATION_DENIED_OBJECT_MANAGED**

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

- **MAPPED_HOSTS_NOT_IN_DESTINATION**

  A host that is mapped to a volume in the pool is not associated with the target domain.

- **MAPPED_CLUSTERS_NOT_IN_DESTINATION**

  A cluster that is mapped to a volume in the pool is not associated with the target domain.

- **DOMAIN_USED_SCHEDULE_NOT_IN_DESTINATION**

  A schedule used by a mirror in the pool is not associated with the target domain.

- **DOMAIN_USED_TARGET_NOT_IN_DESTINATION**

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

- **DOMAIN_MAX_MIRRORS_REACHED**

  The domain exceeds the maximum allowed number of mirrors.

- **DOMAIN_MAX_DMS_REACHED**

  The domain exceeds the maximum allowed number of data migrations.

- **DOMAIN_MAX_VOLUMES_REACHED**

  The domain exceeds the maximum allowed number of volumes.

- **MAX_MIRRORS_REACHED**

  The maximum number of mirrors is already reached.

- **MAX_DMS_REACHED**

  The maximum number of remote volumes (mirror/migration) is already reached.

  **Troubleshooting:** Delete unnecessary data migration objects.

- **MAX_VOLUMES_REACHED**

  The maximum allowed number of volumes is already reached.

- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support

# 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.
- **snmp_trap_type** - The valid options are: common, severity_specific.
- **snmp_type** - SNMP version the snmpd should work with. The valid options are: none, v2c, v3, both.
- **support_center_port_type** - The valid options are: Management, VPN.
- **system_id** - Unique system identifier (equivalent to a serial number).

© Copyright IBM Corp. 2016, 2022

- **system_name**
- **snmpv3_user**
- **snmpv3_encryption_type** - The valid options are: DES, AES.
- **snmpv3_encryption_passphrase**
- **snmpv3_authentication_type** - The valid options are: MD5, SHA.
- **snmpv3_authentication_passphrase**
- **ipv6_state** - Is IPv6 support enabled.
- **ipsec_state** - Is IPSec support enabled.
- **ipsec_track_tunnels** - Should notify when IPSec tunnels are opened and closed.
- **impending_power_loss_detection_method** - The valid options are: BBU, None.
- **async_min_rpo** - Shortest ASync RPO allowed in the system.
- **async_max_rpo** - Longest ASync RPO allowed in the system.
- **reserved_free_physical_capacity** - Reserved physical space for OOPS recovery, in XAU units.

**Example:**

```
config_get
```

**Output:**

```
Name                         Value
------------------------ ----------------------------------
email_sender_address     support@example.com
email_reply_to_address   storage@example.com
dns_primary              192.0.2.1
dns_secondary
iscsi_name               iqn.2005-10.com.xivstorage:010140
system_name              IBM Storage System
```

## Access control

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

## Return codes

- **CONF_SERVER_UNREACHABLE**

  The configuration server is unreachable.
- **UNRECOGNIZED_CONFIG_PARAMETER**

  Unrecognized configuration parameter: '*name*'.

  **Troubleshooting:** Use a valid configuration parameter in the command syntax. For the list of valid configuration parameters, see the CLI Reference Guide.

# Setting configuration parameters

Use the **config_set** command to set configuration parameters.

```
config_set name=Name value=ParamValue
```

## Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **name** | String | Name of the parameter to set. | Y |
| **value** | String | Value of the parameter. | Y |

This command sets the values of configuration parameters.

The values of the following parameters can be set:

- **dns_master** - IP address of the master DNS server.
- **dns_slave** - IP address of the slave DNS server.
- **email_sender_address** - Email address used as the sender's address when sending email messages. Once set, this parameter cannot be set to null.
- **email_reply_to_address** - Reply-to address to be used when sending emails. This is useful for troubleshooting errors in email addresses.
- **system_name** - Name used as the sender's name when sending email messages.
- **defaultuser** - Default user to be used if no user is specified for the CLI. If null, a user must be specified.
- **snmp_sysname** - SNMP system name as shown in the SNMP MIB. (.1.3.6.1.2.1.1.5.0)
- **snmp_location** - SNMP location as shown in the SNMP MIB. (.1.3.6.1.2.1.1.6.0)
- **snmp_contact** - SNMP contact as shown in the SNMP MIB. (.1.3.6.1.2.1.1.4.0 )
- **email_subject_format** - Controls the formatting of the email subject line. To insert the event's data, use the following tags: **{severity}**, **{description}**, or **{system_name}**. System default is **"{severity}: {description}"**.
- **ntp_server** - IP address or DNS name of the NTP server.
- **snmp_community** - Community used for SNMP queries of the system.
- **snmp_trap_community** - Community used for SNMP traps sent by the system.
- **snmp_trap_type** - The valid options are: common, severity_specific.
- **snmp_type** - SNMP version the snmpd should work with. The valid options are: none, v2c, v3, both.

**Example:**

```
config_set name=dns_secondary value=192.0.2.2
```

**Output:**

```
Command completed successfully
```

## Access control

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

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

### Return codes

- **UNRECOGNIZED_CONFIG_PARAMETER**

  Unrecognized configuration parameter: '*name*'.

  **Troubleshooting:** Use a valid configuration parameter in the command syntax. For the list of valid configuration parameters, see the CLI Reference Guide.
- **READ_ONLY_CONFIG_PARAMETER**

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

  **Troubleshooting:** Read-only parameters are not available for modifying.
- **IPV4_NOT_CONFIGURED**

  The IPv4 address is not configured on the management interface.

  **Troubleshooting:** Define an IPv4 address for management before disabling IPv6.
- **RULE_WITH_SNMP_DEST_EXISTS**

  Cannot set snmp_type to None. There is a rule that contains an SNMP destination.

## Testing the DNS

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

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

### Parameters

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

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

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

The result of each defined DNS server is displayed.

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

**Example:**

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

**Output:**

```
Name         IP (Primary DNS)   IP (Secondary DNS)
------------ ------------------ --------------------
hermes.xiv   203.0.113.1        Not Found
```

### Access control

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

### Return codes

- **DNS_SERVER_NOT_DEFINED**

  No DNS servers are defined.

## Displaying help

Use the **help** command to display system help.

```
help [ category=Category | search=SearchString | command=CommandName ]
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **category** | String | Category name. | N |
| **search** | String | Search string. | N |
| **command** | String | Command name. | N |

This command displays the help as follows:

- No parameters - Lists all the commands with their short descriptions, grouped by categories.
- Category - Lists all the commands in the category, with their short descriptions.
- Search - Lists the short descriptions of all the commands in which the search string appears in their name or short description.
- Command - Displays the command name and short discrption.

**Example:**

```
help category=volume
```

**Output:**

```
Category  Name       Description
volume    vol_copy   Copies a source volume onto a target volume.
volume    vol_create Creates a new volume.
volume    vol_delete Deletes a volume
volume    vol_format Formats a volume.
volume    vol_list   Lists all volumes, or a specific one.
volume    vol_lock   Locks a volume, so that it is read-only.
volume    vol_rename Renames a volume
volume    vol_resize Resizes a volume
volume    vol_unlock Unlocks a volume, so that it is no longer read-only,
                     and can be written to.
```

| Field ID | Field output | Default position |
|---|---|---|
| **category** | Category | 1 |
| **name** | Name | 2 |
| **access_control** | Access Control | N/A |
| **syntax** | Syntax | N/A |
| **fields** | Fields | N/A |
| **description** | Description | 3 |
| **example** | Example | N/A |

### Access control

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

## Printing the current maintenance urgency

Use the **maintenance_urgency_list** command to display the current maintenance urgency of the system.

```
maintenance_urgency_list
```

**Example:**

```
maintenance_urgency_list
```

**Output:**

```
maintenance_urgency = "NONE"
```

### Access control

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

## Shutting down the system

Use the **shutdown** command to shut down the system.

```
shutdown [ emergency=<yes|no> ] [ ignore_ha=<yes|no> ]
```

### Parameters

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

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

**Attention:**

**USING THIS OPTION MAY CAUSE DATA LOSS.**

#### Example:

```
shutdown -y
```

#### Output:

```
Command completed successfully.
```

### Access control

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

| User Category | Permission |
|---|---|
| Technicians | Disallowed |

### Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_SHUT_DOWN**

  Are you sure you want to shut down the machine and all its components?

### Return codes

- **COMMAND_IS_NOT_VALID_IN_CURRENT_SYSTEM_STATE**

  The requested command cannot be invoked in the current system state.
- **FIRMWARE_UPGRADE_IN_PROGRESS**

  Firmware upgrade in progress.

  **Troubleshooting:** Contact IBM Support.
- **FLASH_CCL_IN_PROGRESS**

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

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

  The system owns volume(s) defined as primary in a HyperSwap relation. If you continue with shutdown without first handling those relations, the host may lose access to those volumes.

  **Troubleshooting:** It is recommended to run switch_roles before continuing. You may explicitly add ignore_ha=yes to force the operation.

## Listing the operational state

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

```
state_list
```

**Example:**

```
state_list
```

**Output:**

```
Category                 Value
------------------------ -----------------
data_protection_status   Fully Protected
data_reduction_state     Online
encryption               Enabled
safe_mode                no
shutdown_reason          No Shutdown
system_state             on
target_state             on
```

### Access control

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

| User Category | Permission |
|---|---|
| Security administrator | Allowed |
| Read-only users | Allowed |
| Technicians | Allowed |

## Displaying system usage and data reduction statistics

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

```
system_usage_get
```

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

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

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

**Example:**

```
system_usage_get
```

**Output:**

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

cont:
Total Written Pending Deletion (GB)   Data Reduction Savings (%)   Total Stored (GB)
------------------------------------ --------------------------- -------------------
0                                    87                          45

cont:
Deduplication Factor   Compression Factor
---------------------- --------------------
1.37                   5.63
```

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

| Field ID | Field output | Default position |
|---|---|---|
| **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 physical and effective capacity

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

```
system_capacity_list
```

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

**Effective capacity**

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

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

The command output for effective capacity includes the following field:

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

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

**Physical capacity**

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

The command output for physical capacity includes the following fields:

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

- **Free (%)** — The free capacity, represented as percentage of the total capacity.

**Example:**

```
system_capacity_list
```

**Output:**

```
Type        Total (GB)   Allocated (GB) Allocated (%)  Free (GB)  Free (%)
----------- ----------- --------------- -------------- ---------- ---------
Effective   N/A         50640           N/A            N/A        N/A
Physical    1653        140             8              1513       92
```

| Field ID | Field output | Default position |
|---|---|---|
| **type** | Type | 1 |
| **total** | Total (GB) | 2 |
| **allocated** | Allocated (GB) | 3 |
| **allocated_percent** | Allocated (%) | 4 |
| **free** | Free (GB) | 5 |
| **free_percent** | Free (%) | 6 |
| **total_MiB** | Total (MiB) | N/A |
| **allocated_MiB** | Allocated (MiB) | N/A |
| **free_MiB** | Free (MiB) | N/A |

### Access control

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

## Displaying information about effective capacity

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

```
system_effective_capacity_get
```

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

**Example:**

```
system_effective_capacity_get
```

**Output:**

```
Limit (GB)   Effective (GB)   Effective Factor
------------ ---------------- ------------------
1400071      N/A              N/A
```

| Field ID | Field output | Default position |
|---|---|---|
| **effective_capacity_limit** | Limit (GB) | 1 |
| **effective_capacity** | Effective (GB) | 2 |
| **effective_capacity_factor** | Effective Factor | 3 |
| **effective_capacity_max_limit** | Max. Limit (GB) | N/A |
| **effective_capacity_default_limit** | Default Limit (GB) | N/A |
| **effective_capacity_min_limit** | Min. Limit (GB) | N/A |

### Access control

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

## Displaying system capacity thresholds

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

```
system_capacity_threshold_list
```

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

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

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

The output includes the following fields: This command displays the following information about the system capacity thresholds:

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

**Example:**

```
system_capacity_threshold_list
```

**Output:**

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

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

### Access control

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

## Changing a system capacity threshold

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

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

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **hysteresis** | Integer | The hysteresis value in per cent (same for all thresholds). | N | N/A |
| **ordinal** | Integer | The ordinal of the threshold. | N | N/A |
| **enabled** | Boolean | Enable or disable system capacity threshold. | N | yes |

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **threshold** | Integer | The new threshold value in percent, strictly monotonically increasing across thresholds. | N | No threshold |
| **severity** | N/A | The new severity value, strictly monotonically increasing across thresholds. | N | No severity |

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

The example below changes the hysteresis of the system capacity.

**Example:**

```
system_capacity_threshold_change hysteresis=3
```

The example below enables ordinal 1 and changes the threshold and severity parameters of the system capacity.

**Example:**

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

### Access control

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

### Return codes

- **SYSTEM_CAPACITY_DUPLICATE_THRESHOLD**

  The new threshold value is a duplicate of another threshold value.
- **SYSTEM_CAPACITY_THRESHOLD_NOT_MONOTONIC**

  The threshold value or severity must be monotone.

## Resuming the system's normal operation

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

```
system_resume_normal_operation
```

This command is used to move a system back to read-write state after it ran out of physical space.

**Example:**

```
system_resume_normal_operation
```

**Output:**

```
Command completed successfully.
```

### Access control

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

### Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_RESUME_NORMAL_OPERATION**

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

### Return codes

- **SYSTEM_IS_STILL_OUT_OF_PHYSICAL_SPACE**

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

  The system is already operating normally.

## Displaying the current time

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

```
time_list
```

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

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

**Example:**

```
time_list
```

**Output:**

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

### Access control

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

## Setting the system's time

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

```
time_set time=Timestamp
```

### Parameters

| Name | Description | Mandatory |
|---|---|---|
| **time** | New current time. | Y |

**Example:**

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

**Output:**

```
Command completed successfully.
```

### Access control

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

### Return codes

- **SYSTEM_TIME_NOT_CHANGED**

  The system time has not changed.

  **Troubleshooting:** Please try again.

- **FLASH_ENCLOSURE_TIME_UPDATE_FAILED**

  Failed to update the flash enclosure time.

  **Troubleshooting:** Please try again.
- **BAD_TIMESTAMP**

  The timestamp cannot be deciphered.

## Listing optional time zones

Use the **timezone_list** command to list all optional time zones.

```
timezone_list
```

Standard POSIX time zones are used. http://www.timeanddate.com/worldclock/ provides a full description of all time zones.

**Example:**

```
timezone_list
```

**Output:**

```
Timezone
---------------------------------
Africa/Abidjan
Africa/Accra

 ...

WET
Zulu
```

| Field ID | Field output | Default position |
|---|---|---|
| **name** | Timezone | 1 |

### Access control

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

## Setting the time zone

Use the **timezone_set** command to set the time zone of the system.

```
timezone_set timezone=TimeZone
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **timezone** | String | New time zone of the system. | Y |

See Listing optional time zones for a complete list of optional time zones.

Standard POSIX time zones are used. http://www.timeanddate.com/worldclock/ provides a full description of all time zones.

**Example:**

```
timezone_set timezone=Etc/GMT+1
```

**Output:**

```
Command completed successfully
```

### Access control

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

### Return codes

- **BAD_TIMEZONE_NAME**

  Timezone is not recognized by the system.

## Printing the current system version

Use the **version_get** command to print the current version of the system.

```
version_get
```

| Field ID | Field output | Default position |
|---|---|---|
| **system_version** | Version | 1 |

**Example:**

```
version_get
```

**Output:**

```
Version
10.2
```

### Access control

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

## Displaying the values of VPD parameters

Use the **vpd_config_get** command to display the values of VPD parameters.

```
vpd_config_get  [ name=Name ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **name** | String | Name of the parameter to print. | N | All parameters. |

| Field ID | Field output | Default position |
|---|---|---|
| **name** | Name | 1 |
| **value** | Value | 2 |

See Setting VPD parameters for a full list of available settings.

**Example:**

```
vpd_config_get name=site.city
```

**Output:**

```
Name        Value
----------- --------
site.city   Gotham
```

### Access control

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

### Return codes

- **CONF_SERVER_UNREACHABLE**

  The configuration server is unreachable.
- **UNRECOGNIZED_CONFIG_PARAMETER**

  Unrecognized configuration parameter: '*name*'.

  **Troubleshooting:** Use a valid configuration parameter in the command syntax. For the list of valid configuration parameters, see the CLI Reference Guide.

## Setting VPD parameters

Use the **vpd_config_set** command to set the values of VPD (Vital Product Data) parameters.

```
vpd_config_set name=Name value=ParamValue
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **name** | String | Name of the parameter to set. | Y |
| **value** | String | Value of the parameter. | Y |

This command sets the following values of VPD parameters, where only the name is mandatory.:

- customer.name
- customer.primary_contact.calling_hours
- customer.primary_contact.email
- customer.primary_contact.mobile_phone
- customer.primary_contact.name
- customer.primary_contact.office_phone
- customer.primary_contact.time_zone
- customer.secondary_contact.calling_hours
- customer.secondary_contact.email
- customer.secondary_contact.mobile_phone
- customer.secondary_contact.name
- customer.secondary_contact.office_phone
- customer.secondary_contact.time_zone
- hardware_info.hw_cable_bundle
- hardware_info.hw_door
- hardware_info.hw_patch_panel
- hardware_info.hw_patch_panel_label
- hardware_info.hw_power_cable_config
- hardware_info.hw_rack_type
- hardware_info.hw_rps
- interface_config.model
- machine_model
- machine_type
- main_ibm_contact.calling_hours

- main_ibm_contact.email
- main_ibm_contact.mobile_phone
- main_ibm_contact.name
- main_ibm_contact.office_phone
- main_ibm_contact.time_zone
- non_mutable_vpd_info.original_flashed_version
- non_mutable_vpd_info.original_flashing_date
- disk_size
- remote_support.customer_contact.calling_hours
- remote_support.customer_contact.email
- remote_support.customer_contact.mobile_phone
- remote_support.customer_contact.name
- remote_support.customer_contact.office_phone
- remote_support.customer_contact.time_zone
- remote_support.modem_phone_number
- remote_support.primary_ibm_ip
- remote_support.secondary_ibm_ip
- remote_support.special_instructions
- remote_support.vpn_ip_1
- remote_support.vpn_ip_2
- site.building_location
- site.city site.country
- site.name
- site.postal_code
- site.state
- site.street_address
- system_info.sys_ec_level
- system_info.sys_hw_level
- system_info.machine_configuration_entry_point
- system_info.PID

**Example:**

```
vpd_config_set name=customer.name value=JohnDoe
```

**Output:**

```
Command completed successfully.
```

## Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Disallowed |

| User Category | Permission |
|---|---|
| Read-only users | Disallowed |
| Technicians | Allowed |

### Return codes

- **READ_ONLY_CONFIG_PARAMETER**

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

  **Troubleshooting:** Read-only parameters are not available for modifying.
- **UNRECOGNIZED_CONFIG_PARAMETER**

  Unrecognized configuration parameter: '*name*'.

  **Troubleshooting:** Use a valid configuration parameter in the command syntax. For the list of valid configuration parameters, see the CLI Reference Guide.

## Displaying the system's MIB file

Use the **mib_get** command to display the system's MIB file.

```
mib_get
```

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

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

**Output:**

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

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

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

| Field ID | Default position |
|---|---|
| **line** | 1 |

**Example:**

```
mib_get
```

**Output:**

```
------------------------------------------
------------------------------------------
------------------------------------------
-- -*- SNMP -*- mode for Emacs
XIV-MIB DEFINITIONS ::= BEGIN

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

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

### Access control

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

### Return codes

- **CANNOT_READ_FROM_FILE**

  Cannot read from file '*Filename*'.

  **Troubleshooting:** Contact IBM Support.

## Retrieving the electronic license acceptance status

Use the **elicense_status_get** command to retrieve the electronic license acceptance status.

```
elicense_status_get
```

**Example:**

```
elicense_status_get
```

**Output:**

```
Status
----------
Accepted
```

| Field ID | Field output | Default position |
|---|---|---|
| **status** | Status | 1 |

### Access control

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

## Retrieving a fragment of the electronic license file

Use the **elicense_blob_get** command to retrieve a fragment of the electronic license file.

```
elicense_blob_get beg=BeginIndex size=Number
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **beg** | Positive integer | Beginning of the fragment in bytes. | Y |
| **size** | Positive integer | Length of the fragment in bytes. The maximum length allowed is 1000000. | Y |

**Example:**

```
elicense_blob_get beg=0 size=20
```

**Output:**

```
<file_size value="1300473"/>
<fragment value="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 IBM Support.

# Accepting the electronic license agreement

Use the **elicense_accept** command to accept the electronic license agreement.

```
elicense_accept version=Version [ approver_name=UserName ]
```

## Parameters

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

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

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

**Example:**

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

## Access control

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

## Return codes

- **ELICENSE_INCOMPATIBLE_VERSION**

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

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

  Electronic license already accepted

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

  Electronic license check is disabled

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

# Enabling command auditing

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

```
audit_enable
```

This command is used by a security administrator to enable the auditing of user-entered CLI commands on an external auditing server. For this command to complete successfully, the current auditing state must be DISABLED (that is, the **audit_show** command returns a *no*), and at least one audit server must be configured successfully by the **audit_config_set** command.

**Note:**

The **audit_enable** command has a default filter which filters out the following commands: **appadmin_capabilities_get**, **domain_global_list**, **state_list**, **metadata_list**, **snap_group_list**, **user_metadata_list**, and **version_get**. For more information about filters and filter changes, contact IBM Support.

**Example:**

```
audit_enable
```

**Output:**

```
Command completed successfully.
```

## Access control

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

## Return codes

- **AUDIT_ALREADY_ENABLED**

  Command auditing is already enabled.
- **AUDIT_NO_AUDIT_SERVER_DEFINED**

  No audit logging server is configured.

# Disabling command auditing

Use the **audit_disable** command to disable CLI command auditing.

```
audit_disable
```

This command disables command auditing, provided that auditing is currently enabled, that is the **audit_show** command returns a *yes*.

**Example:**

```
audit_disable -y
```

**Output:**

```
Command completed successfully.
```

### Access control

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

### Return codes

- **AUDIT_NOT_ENABLED**

  Command auditing is not enabled.

## Displaying the command audit state

Use the **audit_show** command to show the current state of CLI command auditing.

```
audit_show
```

| Field ID | Field output | Default position |
|---|---|---|
| **audit_enabled** | Auditing Enabled | 1 |

**Example:**

```
audit_show
```

**Output:**

```
Auditing Enabled
------------------
yes
```

### Access control

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

# Configuring audit servers

Use the **audit_config_set** command to configure CLI command auditing.

```
audit_config_set primary_server=Address [ primary_port=port ] [ secondary_server=Address ]
 [ secondary_port=port ] [ protocol=protocol ]
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **primary_server** | N/A | IP address of the primary auditing server. | Y | N/A |
| **primary_port** | Positive integer | IP port number of the primary auditing server. | N | Default for protocol |
| **secondary_server** | N/A | IP address of the secondary auditing server. | N | empty |
| **secondary_port** | Positive integer | IP port number of the secondary auditing server. | N | Default for protocol |
| **protocol** | Enumeration | Transport protocol. Only RFC-5424 Syslog over UDP is currently supported. | N | syslog |

This command configures the primary and, optionally, the secondary auditing server for CLI command logging.

## Access control

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

# Checking the command audit state

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

```
audit_config_get
```

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

**Example:**

```
audit_config_get
```

**Output:**

```
Primary Server   Primary Port   Secondary Server   Secondary Port   Protocol
--------------- -------------- ------------------ ---------------- ----------
192.0.2.1        514                               0                syslog
```

### Access control

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

## Retrieving the list of Flash VDisks

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

```
flash_vdisk_list
```

**Example:**

```
flash_vdisk_list
```

**Output:**

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

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

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

### Access control

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

## Enabling CIM service

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

```
cim_enable
```

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

**Example:**

```
cim_enable
```

**Output:**

```
Command completed successfully.
```

### Access control

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

### Return codes

- **CIM_ALREADY_ENABLED**

  The CIM port is already enabled.

## Disabling the CIM service

Use the **cim_disable** command to disable the CIM service.

```
cim_disable
```

This command disables the CIM service and the associated SLP service. In order for this command to complete successfully, the current CIM service state must be ENABLED (that is, the **cim_show** command returns *yes*).

**Example:**

```
cim_disable
```

**Output:**

```
Command completed successfully.
```

### Access control

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

### Return codes

- **CIM_NOT_ENABLED**

  The CIM port is not enabled.

## Displaying the CIM service state

Use the **cim_show** command to display the current state of CIM service.

```
cim_show
```

| Field ID | Field output | Default position |
|---|---|---|
| **cim_enabled** | CIM Enabled | 1 |

**Example:**

```
cim_show
```

**Output:**

```
CIM Enabled
------------------
yes
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Allowed |

| User Category | Permission |
| --- | --- |
| 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**

  The target name does not exist.
- **TARGET_INVALID_CONNECTION_DURATION_THRESHOLD**

  The target connection duration threshold should be in the [1,1000000] range.

© Copyright IBM Corp. 2016, 2022

# Updating the target's mirroring configuration

Use the **target_config_sync_rates** command to change the target's mirroring configuration.

```
target_config_sync_rates target=TargetName [ max_initialization_rate=MaxInitializationRate ]
[ max_syncjob_rate=MaxSyncjobRate ] [ max_resync_rate=MaxResyncRate ]
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **target** | Object name | The updated target. | Y | N/A |
| **max_initialization_rate** | Positive integer | Specifies the maximum rate for initial synchronization. Cannot be larger than **max_syncjob_rate**. | N | Unchanged |
| **max_syncjob_rate** | Positive integer | Specifies the default maximum rate for sync job synchronization. Cannot be larger than **max_resync_rate**. | N | Unchanged |
| **max_resync_rate** | Positive integer | Specifies the maximum rate for re-synchronization | N | Unchanged |

This command changes the system ID of the remote target. The synchronization rate units are MB per second. The default rates are: 100 MB/s for initialization rate, 300 MB/s for resync rate. The default system_id is the value that is set with the **config_set** command.

## Access control

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

## Return codes

- **TARGET_BAD_NAME**

  The target name does not exist.

- **TARGET_ILLEGAL_RATE_VALUES**

  The maximum initialization rate should be smaller than or equal to the maximum synchronization job rate. The maximum synchronization job rate should not be greater than the maximum resynchronization rate.

# Activating connectivity to a remote target

Use the **target_connectivity_activate** command to activate connectivity between a port on the local storage system and a port on a remote target.

```
target_connectivity_activate target=TargetName
      < ipaddress=IPaddress local_ipinterface=IPInterface > |
<
      fcaddress=wwpn local_port=PortID >
```

## Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **target** | Object name | Remote target of the connectivity definition. | Y |
| **ipaddress** | N/A | IP address of the port on the remote target (iSCSI targets only). | N |
| **local_ipinterface** | Object name | Local IP interface to be connected to the remote port (iSCSI only) | N |
| **fcaddress** | N/A | FC address of the port on the remote target (FC targets only). | N |
| **local_port** | N/A | Port identifier. | N |

Each connectivity definition can be either active or inactive. The system does not use inactive connectivity definitions. Target connectivity is active by default.

This command has no effect if the connectivity is already active.

## Access control

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

## Return codes

- **TARGET_BAD_NAME**

  The target name does not exist.
- **CONNECTIVITY_NOT_DEFINED**

  No remote port is connected through this local port.
- **COMPONENT_IS_NOT_AN_FC_PORT**

  An FC port must be specified for the component.
- **COMMAND_NOT_ALLOWED_ON_MANAGEMENT_OR_VPN_INTERFACE**

  The operation is not allowed on the management or VPN IP Interface.
- **IPINTERFACE_DOES_NOT_EXIST**

This IP interface name does not exist.

- **TARGET_PORT_BAD_ADDRESS**

  The remote port address is illegal or does not belong to the remote target.

- **BAD_LOCAL_IP_PORT**

  The ID of a local IP port must be specified.

## Deactivating connectivity to a remote target

Use the **target_connectivity_deactivate** command to deactivate connectivity between a port on the local storage system and a port on a remote target.

```
target_connectivity_deactivate target=TargetName
< ipaddress=IPaddress local_ipinterface=IPInterface > |
<
fcaddress=wwpn  local_port=PortID > [ force_on_olvm_peer=<yes|no> ]
 [ force_on_ha_peer=<yes|no> ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **target** | Object name | Remote target of the connectivity definition. | Y | N/A |
| **ipaddress** | N/A | IP address of the port on the remote target (iSCSI targets only). | N | N/A |
| **local_ipinterface** | Object name | Local IP interface that is connected to the remote port (iSCSI only). | N | N/A |
| **fcaddress** | N/A | FC address of the port on the remote target (FC targets only). | N | N/A |
| **local_port** | N/A | Port identifier. | N | N/A |
| **force_on_olvm_peer** | Boolean | Informs the system whether the command should be applied on an olvm peer. | N | No |
| **force_on_ha_peer** | Boolean | Force the deactivation on a HyperSwap target. | N | No |

This command deactivates connectivity.

Each connectivity definition can be either active or inactive. The system does not use inactive connectivity definitions. Target connectivity is active by default. Connectivity can be reactivated using Activating connectivity to a remote target.

This command has no effect if the connectivity is already deactivated.

**Example:**

```
target_connectivity_deactivate
target=Nextra2 local_module=101
```

**Output:**

```
Command completed successfully
```

### Access control

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

### Return codes

- **TARGET_BAD_NAME**

  The target name does not exist.
- **COMMAND_NOT_ALLOWED_ON_MANAGEMENT_OR_VPN_INTERFACE**

  The operation is not allowed on the management or VPN IP Interface.
- **CONNECTIVITY_NOT_DEFINED**

  No remote port is connected through this local port.
- **COMPONENT_IS_NOT_AN_FC_PORT**

  An FC port must be specified for the component.
- **TARGET_PORT_BAD_ADDRESS**

  The remote port address is illegal or does not belong to the remote target.
- **BAD_LOCAL_IP_PORT**

  The ID of a local IP port must be specified.
- **IPINTERFACE_DOES_NOT_EXIST**

  This IP interface name does not exist.
- **TARGET_HAS_OLVM_RELATIONSHIP**

  The target has an IBM Hyper-Scale Mobility relationship, and therefore cannot be deactivated or deleted.
- **TARGET_HAS_HA_RELATIONSHIP**

  The target has an IBM HyperSwap relationship, and therefore cannot be deactivated or deleted.

## Defining connectivity to a remote target

Use the **target_connectivity_define** command to define connectivity between a port on the local storage system and a port on a remote target.

```
target_connectivity_define target=TargetName < ipaddress=IPaddress local_ipinterface=IPInterface
> | < fcaddre
ss=wwpn local_port=PortID >
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **target** | Object name | Remote target of the connectivity definition. | Y |

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **ipaddress** | N/A | IP address of the port on the remote target (iSCSI targets only). | N |
| **local_ipinterface** | Object name | Local IP interface to be connected to the remote port (iSCSI only). | N |
| **fcaddress** | N/A | FC address of the port on the remote target (FC targets only). | N |
| **local_port** | N/A | FC port (FC only). | N |

Connectivity between a local and a target storage system is defined between a specific port on a local storage system and a port on the target storage system.

Each connectivity definition can be either active or inactive. The system does not use inactive connectivity definitions. Target connectivity is active by default. An option is provided to deactivate (**target_connectivity_deactivate**) and then reactivate (**target_connectivity_activate**) it, if required. Target connectivity can be deleted (Deleting connectivity to a remote target) and a list of target connectivity definitions (Listing target connectivity definitions) can be displayed.

If the port is used for target connectivity, then only one IP interface is supported. In this case, the other IP interfaces must be removed from the Ethernet port in order to define target connectivity.

## Access control

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

## Return codes

- **TARGET_BAD_NAME**

  The target name does not exist.
- **CONN_EXISTS**

  A remote port is already connected through this local port.
- **MAX_CONNECTIONS_REACHED**

  The maximum number of connections is already reached.
- **MAX_ISCSI_CONNECTIONS_PER_MODULE_REACHED**

  The maximum number of iSCSI connectivities for that module is already reached.
- **COMPONENT_IS_NOT_AN_FC_PORT**

  An FC port must be specified for the component.
- **COMPONENT_IS_NOT_FC_INITIATOR_PORT**

  An FC initiator port must be specified for the component.
- **BAD_LOCAL_IP_PORT**

  The ID of a local IP port must be specified.

- **COMMAND_NOT_ALLOWED_ON_MANAGEMENT_OR_VPN_INTERFACE**

  The operation is not allowed on the management or VPN IP Interface.
- **MORE_THAN_ONE_IP_INTERFACE_DEFINED_ON_PORT**

  There is more than one IP Interface defined on port.
- **IPINTERFACE_DOES_NOT_EXIST**

  This IP interface name does not exist.
- **TARGET_PORT_BAD_ADDRESS**

  The remote port address is illegal or does not belong to the remote target.

## Deleting connectivity to a remote target

Use the **target_connectivity_delete** command to delete connectivity between a port on the local storage system and a port on a remote target.

```
target_connectivity_delete target=TargetName < ipaddress=IPaddress local_ipinterface=IPInterface
> | < fcaddre
ss=wwpn local_port=PortID > [ force_on_olvm_peer=<yes|no> ] [ force_on_ha_peer=<yes|no> ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **target** | Object name | Remote target of the connectivity definition. | Y | N/A |
| **ipaddress** | N/A | IP address of the port on the remote target (iSCSI targets only). | N | N/A |
| **local_ipinterface** | Object name | Local IP interface that is connected to the remote port (iSCSI only). | N | N/A |
| **fcaddress** | N/A | FC address of the port on the remote target (FC targets only). | N | N/A |
| **local_port** | N/A | Port number on the local module (FC only). | N | N/A |
| **force_on_olvm_peer** | Boolean | Informs the system whether the command should be applied on an IBM Hyper-Scale Mobility peer. | N | No |
| **force_on_ha_peer** | Boolean | Force the connectivity deletion on a HyperSwap target. | N | No |

Only a previously defined connectivity definition can be deleted.

**Example:**

```
target_connectivity_delete target=XIV2 local_module=101
```

**Output:**

```
Command completed successfully
```

### Access control

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

### Return codes

- **TARGET_BAD_NAME**

  The target name does not exist.
- **COMMAND_NOT_ALLOWED_ON_MANAGEMENT_OR_VPN_INTERFACE**

  The operation is not allowed on the management or VPN IP Interface.
- **CONNECTIVITY_NOT_DEFINED**

  No remote port is connected through this local port.
- **COMPONENT_IS_NOT_AN_FC_PORT**

  An FC port must be specified for the component.
- **TARGET_PORT_BAD_ADDRESS**

  The remote port address is illegal or does not belong to the remote target.
- **BAD_LOCAL_IP_PORT**

  The ID of a local IP port must be specified.
- **IPINTERFACE_DOES_NOT_EXIST**

  This IP interface name does not exist.
- **TARGET_HAS_OLVM_RELATIONSHIP**

  The target has an IBM Hyper-Scale Mobility relationship, and therefore cannot be deactivated or deleted.
- **TARGET_HAS_HA_RELATIONSHIP**

  The target has an IBM HyperSwap relationship, and therefore cannot be deactivated or deleted.

## Listing target connectivity definitions

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

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

### Parameters

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

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

### Access control

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

## Defining a remote target

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

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

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **target** | Object name | Local name of the remote target. | Y | N/A |
| **protocol** | Enumeration | FC (Fiber Channel) or iSCSI, depending on the communication protocol supported by the remote host. | Y | N/A |
| **iscsi_name** | iSCSI initiator name | iSCSI name of the remote target. This field is mandatory for iSCSI hosts. | N | N/A |
| **system_id** | String | ID of the remote system. Should be the same as the output of Displaying the values of configuration parameters of the *system_id* variable on the remote system. | N | N/A |
| **xiv_features** | Boolean | Defines the remote system as an XIV system. Non-XIV systems are used only for data migration. | N | Yes |

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **domain** | N/A | The cluster will be attached to the specified domains. To define more than one domain, separate them with a comma. To specify all existing domains, use "*". | N | none |
| **quorum_witness** | Object name | The name of the Quorum Witness that is associated with the target. | N | none |
| **uses_512b_sectors** | Boolean | Optimize the Asynchronous mirror data transfer for remote targets with 512B sector size. | N | No |

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

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

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

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

**Example:**

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

**Output:**

```
Command completed successfully.
```

## Access control

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

## Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_DEFINE_ANOTHER_TARGET_ON_SYSTEM**

  Defining more than one target to the same remote system is not supported, and may compromise the data on the secondary system. Are you sure the remote system is not already defined as a target?

### Return codes

- **MAX_TARGETS_REACHED**

  The maximum number of defined targets is already reached.
- **TARGET_NAME_EXISTS**

  The target name is already assigned to another target.
- **TARGET_ISCSI_MUST_HAVE_A_NAME**

  iSCSI Target must have an iscsi_name.
- **ISCSI_NAME_NOT_ALLOWED_FOR_FC**

  The FC Target does not have an iscsi_name.
- **TARGET_BAD_SCSI_TYPE**

  The target SCSI type does not exist.
- **DOMAIN_DOESNT_EXIST**

  The domain does not exist.
- **QUORUM_WITNESS_BAD_NAME**

  The Quorum Witness name does not exist.
- **QUORUM_WITNESS_IS_NOT_ACTIVATED**

  The Quorum Witness is not activated.
- **QUORUM_WITNESS_CANNOT_BE_ADDED_TO_A_TARGET_OF_THIS_TYPE**

  A Quorum Witness cannot be added to either iSCSI or non-Spectrum Accelerate target.

## Deleting a remote target

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

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

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **target** | Object name | Target that is deleted. | Y | N/A |
| **force_on_olvm_peer** | Boolean | Informs the system whether the command should be applied on an IBM Hyper-Scale Mobility peer. | N | No |

A target that contains port definitions cannot be deleted. A target with remote mirroring or data migration definitions cannot be deleted.

**Example:**

```
target_delete target=Nextra2
```

**Output:**

```
Command completed successfully.
```

### Access control

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

### Return codes

- **TARGET_BAD_NAME**

  The target name does not exist.
- **TARGET_HAS_PORTS**

  There are ports defined for this target.
- **TARGET_HAS_ASSOCIATIONS**

  There are remote volumes defined on this target.
- **TARGET_HAS_OLVM_RELATIONSHIP**

  The target has an IBM Hyper-Scale Mobility relationship, and therefore cannot be deactivated or deleted.
- **TARGET_HAS_HA_RELATIONSHIP**

  The target has an IBM HyperSwap relationship, and therefore cannot be deactivated or deleted.

## Listing remote targets

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

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

### Parameters

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

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

| Field ID | Field output | Default position |
|---|---|---|
| **name** | Name | 1 |
| **scsi_type** | SCSI Type | 2 |
| **connected** | Mirror Connectivity | 3 |
| **ha_connected** | HA Connectivity | 4 |
| **max_initialization_rate** | Max Initialization Rate | 5 |
| **max_resync_rate** | Max Resync Rate | 6 |
| **max_syncjob_rate** | Max Syncjob Rate | 7 |

| Field ID | Field output | Default position |
|---|---|---|
| **machine_serial_number** | Target Serial Number | 8 |
| **system_id** | System ID | N/A |
| **quorum_witness** | Quorum Witness | 9 |
| **xiv_target** | XIV Target | N/A |
| **iscsi_name** | iSCSI Name | N/A |
| **num_ports** | Number of Ports | N/A |
| **creator** | Creator | N/A |
| **connectivity_lost_event_threshold** | Connection Threshold | N/A |
| **peer_health** | Peer Health | N/A |
| **peer_health_reason** | Peer Health Reason | N/A |
| **peer_qw_configuration** | Peer QW Configuration | N/A |
| **coordinated_qw_lapse** | Coordinated QW Lapse | N/A |
| **arch** | Remote Arch | N/A |

### Access control

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

## Allowing remote mirroring access

Use the **target_mirroring_allow** command to allow remote mirroring operations initiated from a remote target.

```
target_mirroring_allow target=TargetName
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **target** | Object name | Remote target name. | Y |

This command is performed on a local storage system in order to allow the target storage system to read, write, view, create volumes and define the existing volumes as slaves. This command is used when allowing remote mirroring operations. Otherwise, the target storage system cannot access the local storage system. This command also allows a remote target to read and write through the SCSI interface.

Once mirroring is allowed, this permission cannot be revoked.

This operation should also be run on the target storage system so that it gives permission to the local storage system to access it.

This step must be performed before mirroring is defined (**mirror_create**).

**Example:**

```
target_mirroring_allow target=Nextra2
```

**Output:**

```
Command completed successfully.
```

### Access control

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

### Return codes

- **TARGET_BAD_NAME**

  The target name does not exist.
- **TARGET_BAD_TYPE**

  The target machine is not an XIV machine.

## Activating a port

Use the **target_port_activate** command to activate a port on a remote target.

```
target_port_activate target=TargetName < ipaddress=IPaddress | fcaddress=wwpn >
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **target** | Object name | Remote target of the port. | Y |
| **ipaddress** | N/A | IP address of the port on the remote target (iSCSI targets only). | N |
| **fcaddress** | N/A | FC address of the port on the remote target (FC targets only). | N |

Each port in a remote system can be configured as either active or inactive. The system does not use inactive ports. After a port is defined, it is active by default. This command reactivates a port if it was de-activated (by using the **target_port_deactivate** command).

This command has no effect, if the port is already active.

**Example:**

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

**Output:**

```
Command completed successfully
```

### Access control

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

### Return codes

- **TARGET_PORT_BAD_ADDRESS**

  The remote port address is illegal or does not belong to the remote target.
- **TARGET_BAD_PORT_STATE**

  The port is already in the requested activation state.
- **TARGET_BAD_NAME**

  The target name does not exist.

## Adding a new port to a remote target

Use the **target_port_add** command to add a port to a remote target.

```
target_port_add target=TargetName < ipaddress=IPaddress | fcaddress=wwpn >
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **target** | Object name | Remote target to which to add the port. | Y |
| **ipaddress** | N/A | IP address of the port on the remote target (for iSCSI type targets only). | N |
| **fcaddress** | N/A | FC address of the remote port (for FC type targets only). | N |

This command adds a new port to a specified target. A port can be either FC or iSCSI, and its type must conform to the remote target's communication protocol type.

Specify the IP address or the FC address according to communication protocol of the target.

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Allowed |
| Application administrator | Disallowed |
| Security administrator | Disallowed |

| User Category | Permission |
|---|---|
| Read-only users | Disallowed |
| Technicians | Disallowed |

### Return codes

- **TARGET_BAD_NAME**

  The target name does not exist.
- **PORT_EXISTS**

  The port is already defined.
- **MAX_PORTS_REACHED**

  The maximum number of ports defined in the system is already reached.
- **TARGET_PORT_BAD_ADDRESS**

  The remote port address is illegal or does not belong to the remote target.
- **ISCSI_HOST_ILLEGAL_PORT_NAME**

  The port name for iSCSI Host is illegal.

  **Troubleshooting:** Port names for iSCSI Hosts must contain only printable characters.
- **HOST_PORT_EXISTS**

  A host with this port ID is already defined.

## Deactivating a port

Use the **target_port_deactivate** command to deactivate a port of a remote target.

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

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **target** | Object name | The remote target that includes the port to be deactivated. | Y | N/A |
| **ipaddress** | N/A | IP address of the port on the remote target (iSCSI targets only). | N | N/A |
| **fcaddress** | N/A | FC address of the port on the remote target (FC targets only). | N | N/A |
| **force_on_olvm_peer** | Boolean | Informs the system whether the command should be applied on an OLVM peer. | N | No |
| **force_on_ha_peer** | Boolean | Force the deactivation on a HyperSwap target. | N | No |

Each port in a remote system can be configured as either active or in-active. The system does not use an inactive port. After a port is defined, it is active by default. To re-activate a port, issue the **target_port_activate** command (see Activating a port).

**Example:**

```
target_port_deactivate target=XIV2 fcaddress=10:00:00:17:38:27:ec:11
```

**Output:**

```
Command completed successfully
```

### Access control

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

### Return codes

- **TARGET_BAD_NAME**

  The target name does not exist.
- **TARGET_PORT_BAD_ADDRESS**

  The remote port address is illegal or does not belong to the remote target.
- **TARGET_BAD_PORT_STATE**

  The port is already in the requested activation state.
- **TARGET_HAS_OLVM_RELATIONSHIP**

  The target has an IBM Hyper-Scale Mobility relationship, and therefore cannot be deactivated or deleted.
- **TARGET_HAS_HA_RELATIONSHIP**

  The target has an IBM HyperSwap relationship, and therefore cannot be deactivated or deleted.

## Deleting a port from a remote system

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

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

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **target** | Object name | Remote target from which the port is that is deleted. | Y | N/A |
| **ipaddress** | N/A | IP address of the port (for iSCSI targets only). | N | N/A |
| **fcaddress** | N/A | FC address of the remote port (for FC targets only). | N | N/A |
| **force_on_ha_peer** | Boolean | Force the deactivation on a HyperSwap target. | N | No |

**Example:**

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

**Output:**

```
Command completed successfully.
```

### Access control

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

### Return codes

- **TARGET_BAD_NAME**

  The target name does not exist.
- **TARGET_PORT_BAD_ADDRESS**

  The remote port address is illegal or does not belong to the remote target.
- **TARGET_PORT_HAS_CONNECTIVITY**

  Connectivity to this port is already defined.
- **TARGET_HAS_OLVM_RELATIONSHIP**

  The target has an IBM Hyper-Scale Mobility relationship, and therefore cannot be deactivated or deleted.
- **TARGET_HAS_HA_RELATIONSHIP**

  The target has an IBM HyperSwap relationship, and therefore cannot be deactivated or deleted.

## Listing the ports of a remote target

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

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

### Parameters

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

| Field ID | Field output | Default position |
|---|---|---|
| **target_name** | Target Name | 1 |
| **scsi_type** | Port Type | 2 |

| Field ID | Field output | Default position |
|---|---|---|
| **active** | Active | 3 |
| **fc_wwpn** | WWPN | 4 |
| **iscsi_ip_addr** | iSCSI Address | 5 |
| **iscsi_ip_port** | iSCSI Port | 6 |

### Access control

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

## Renaming a remote target

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

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

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **target** | Object name | The target to be renamed. | Y |
| **new_name** | Object name | New name of the target. | Y |

**Example:**

```
target_rename target=Nextra2 new_name=Nextra-DRP
```

**Output:**

```
Command completed successfully.
```

### Access control

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

### Return codes

- **TARGET_BAD_NAME**

The target name does not exist.

- **TARGET_NAME_EXISTS**

  The target name is already assigned to another target.

## Updating the target configuration

Use the **target_update** command to update the target's configuration.

```
target_update target=TargetName [ system_id=SystemId ] [ uses_512b_sectors=<yes|no> ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **target** | Object name | Target to be updated. | Y | N/A |
| **system_id** | String | ID of the remote system. Should be the same as the output of Displaying the values of configuration parameters of the *system_id* variable on the remote system. | N | none |
| **uses_512b_sectors** | Boolean | Optimize the Asynchronous mirror data transfer for remote targets with 512B sector size. | N | No |

This command changes the system ID or other attributes of the remote target.

### Access control

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

### Return codes

- **TARGET_BAD_NAME**

  The target name does not exist.

- **TARGET_HAS_QUORUM_WITNESS_UPDATE_NOT_ALLOWED**

  Updating the target's system ID is not allowed when there is a Quorum Witness defined on the target.

## Adding a Quorum Witness to a target

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

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

### Parameters

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

**Example:**

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

**Output:**

```
Command completed successfully
```

### Access control

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

### Return codes

- **TARGET_BAD_NAME**

  The target name does not exist.
- **QUORUM_WITNESS_BAD_NAME**

  The Quorum Witness name does not exist.
- **TARGET_HAS_A_QUORUM_WITNESS**

  The target already has a Quorum Witness.
- **QUORUM_WITNESS_IS_NOT_ACTIVATED**

  The Quorum Witness is not activated.
- **QUORUM_WITNESS_CANNOT_BE_ADDED_TO_A_TARGET_OF_THIS_TYPE**

  A Quorum Witness cannot be added to either iSCSI or non-Spectrum Accelerate target.

## Removing a Quorum Witness from a target

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

```
target_remove_quorum_witness target=TargetName
```

## Parameters

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

**Example:**

```
target_remove_quorum_witness target=t1
```

**Output:**

```
Command completed successfully
```

## Access control

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

## Return codes

- **TARGET_BAD_NAME**

  The target name does not exist.
- **TARGET_HAS_NO_QUORUM_WITNESS**

  The local target does not have a Quorum Witness defined.
- **TARGET_HAS_ENABLED_HA**

  There are HyperSwap relations with enabled automatic failover configured with this target. Prior to changing the Quorum Witness, the automatic failover must be disabled.

# 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 cancelled. | N | N/A |
| **cg** | Object name | Name of the (local) master consistency group whose non-started snapshot mirrors should be cancelled. | N | N/A |
| **target** | Object name | Target mirror name. Mandatory if 2 mirrors are defined on the volume. | N | [none] |

Only sync jobs that have not started are cancelled. The command does not delete the snapshots themselves.

Upon running the command:

- A warning message is presented to the user for confirmation.
- An event is generated.
- Non-started snapshot mirrors are cancelled.

The command fails under the following conditions:

- The command is issued on a slave volume or consistency group.

### Access control

| User Category | Permission | Condition |
|---|---|---|
| Storage administrator | Allowed | N/A |
| 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 |

© Copyright IBM Corp. 2016, 2022

| User Category | Permission | Condition |
|---|---|---|
| Read-only users | Disallowed | N/A |
| Technicians | Disallowed | N/A |

## Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_CANCEL_SNAPSHOT_MIRRORS_FOR_THE_VOLUME**

  Are you sure you want to delete snapshot mirrors for *Volume*?
- **ARE_YOU_SURE_YOU_WANT_TO_CANCEL_SNAPSHOT_MIRRORS_FOR_THE_CONSISTENCY_GROUP**

  Are you sure you want to delete snapshot mirrors for *Consistency Group*?

## Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **VOLUME_NO_MIRROR**

  The local volume does not have remote mirroring definitions.
- **CONS_GROUP_BAD_NAME**

  The consistency group name does not exist.
- **CONS_GROUP_NO_MIRROR**

  The local consistency group does not have remote mirroring definitions.
- **LOCAL_PEER_IS_NOT_MASTER**

  The local peer is not primary.
- **VOLUME_BELONGS_TO_MIRRORED_CONS_GROUP**

  The volume mirror is part of a consistency group mirror.
- **MIRROR_RETRY_OPERATION**

  There is an operation in progress on this mirror.

  **Troubleshooting:** Retry the command in a few seconds.
- **TARGET_BAD_NAME**

  The target name does not exist.
- **VOLUME_HAS_MULTIPLE_MIRRORS**

  The volume has multiple mirrors. The operation is not allowed or a target must be specified.
- **CONS_GROUP_HAS_MULTIPLE_MIRRORS**

  The consistency group has multiple mirrors. The operation is not allowed, or a target must be specified.
- **VOLUME_TARGET_MISMATCH**

  The volume and target do not match.
- **CONS_GROUP_BAD_TARGET**

  The target name does not match the consistency group.
- **MIRROR_IS_STANDBY**

  The mirror is marked as Standby.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support

# Creating a snapshot mirror (ad hoc sync job)

Use the **mirror_create_snapshot** command to create a snapshot mirror.

```
mirror_create_snapshot <vol=VolName | cg=cgName> [ target=TargetName ]
 < <name=Name [ delete_priority=del_value ]
 > | overwrite=Name> < < slave_name=SnapshotName [ slave_delete_priority=del_value ]
 > | slave_overwrite=SnapshotName>
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **vol** | Object name | The name of the volume to create a snapshot for. | N | N/A |
| **cg** | Object name | Local master consistency group name. | N | N/A |
| **target** | Object name | Target mirror name. Mandatory if 2 mirrors are defined on the volume. | N | [none] |
| **name** | Object name | The name of the new snapshot. | N | N/A |
| **overwrite** | Object name | The name of an existing snapshot that will be overwritten. | N | N/A |
| **slave_overwrite** | Object name | Name of existing snapshot on slave to overwrite. | N | N/A |
| **delete_priority** | Integer | The deletion priority of the volume's snapshot. | N | 1 |
| **slave_name** | Object name | The name of the new snapshot on the slave. | N | N/A |
| **slave_delete_priority** | Integer | The deletion priority of the slave volume's snapshot. | N | 1 |

In synchronous replication, this command takes a snapshot of the source peer (master) and the target peer (slave) at exactly the same time.

In asynchronous replication, the command establishes a process that takes a point-in-time snapshot of the source peer (master) and synchronizes that point-in-time with the slave. The process sets a new sync job to copy the differences between that snapshot and the most recent snapshot that is guaranteed to be synchronized with the target peer.

Prerequisite (for both synchronous and asynchronous mirroring):

- The coupling has to be operational.

Multiple snapshot mirrors:

- Multiple snapshot mirrors can be issued; each mandates the creation of a corresponding sync job.
- Corresponding sync jobs are queued one after another.

Prioritization of sync jobs:

- The snapshot mirror delays the execution of an interval-based mirror if it is running upon arrival of a new interval.

- The snapshot mirror does not, however, cancel the creation of the interval-based sync job. The interval-based mirror will be calculated based on the differences between the most recent snapshot and the last snapshot mirror.

Precedence of the last snapshot mirror over the last replicated snapshot:

- The last_replicated snapshot of the master will be updated to reflect the completed snapshot mirror. Following the completion of the snapshot mirror, its snapshot is duplicated and the duplicate is named last_replicated (the previous last_replicated snapshot is deleted).

Canceling a snapshot mirror:

- The administrator has the ability to cancel snapshot mirrors that have not yet started.

| **Important:** |
| --- |
| The snapshots created concurrently on the master and slave are identical. |

The snapshot mirror results with two last_replicated snapshots that are different and denoted "Master" and "Slave" accordingly:

- On the slave, a snapshot is taken and named **last_replicated**
- On the master, the pertinent snapshot that is mirrored onto the slave is also named **last_replicated**

The outcome for the synchronous mirroring:

- The master blocks host I/O for the duration of creating the snapshots
- The master completes synchronizing pending writes
- A snapshot of the master and slave is taken
- The master no longer blocks host I/O
- An event is generated

**Using the overwrite and slave_overwrite parameters:**

It is possible to overwrite an existing snapshot or snapshot group either on the Master, the Slave, or both.

To specify a local snapshot or snapshot group to be overwritten, use the **overwrite** parameter. Use the **slave_overwrite** parameter to specify a remote snapshot or snapshot group to be overwritten.

The overwrite and slave_overwrite parameters cause the current content of the volume or consistency group to be copied into an existing snapshot or snapshot groups (indicated as the parameter's argument). The overwritten snapshot or snapshots retain the same SCSI device WWN and the same mapping. As a result, the hosts maintain a continuous mapping of the snapshots, and a rescan or similar operation is not needed. The overwritten snapshot or snapshot group must be an existing snapshot or snapshot group of the respective volume or consistency group.

## Access control

| User Category | Permission | Condition |
| --- | --- | --- |
| Storage administrator | Allowed | N/A |
| Storage integration administrator | Allowed | N/A |
| Application administrator | Conditionally Allowed | The volume is mapped to a host or a cluster associated with the user. If a snapshot overwrite is used, the target snapshot must be one created by a server administrator. |
| Security administrator | Disallowed | N/A |
| Read-only users | Disallowed | N/A |
| Technicians | Disallowed | N/A |

## Return codes

- **CONS_GROUP_MISMATCH**

  The snapshot group does not match the consistency group volumes.
- **CONS_GROUP_EMPTY**

  The operation is not allowed on an empty consistency group.
- **CONS_GROUP_BAD_NAME**

  The consistency group name does not exist.
- **CONS_GROUP_NO_MIRROR**

  The local consistency group does not have remote mirroring definitions.
- **LOCAL_PEER_IS_NOT_MASTER**

  The local peer is not primary.
- **MIRROR_IS_NOT_SYNCHRONIZED**

  The mirror is not synchronized.
- **MIRROR_RETRY_OPERATION**

  There is an operation in progress on this mirror.

  **Troubleshooting:** Retry the command in a few seconds.
- **MIRROR_IS_NON_OPERATIONAL**

  The mirror is non-operational.
- **MAX_VOLUMES_REACHED**

  The maximum allowed number of volumes is already reached.
- **DOMAIN_MAX_VOLUMES_REACHED**

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

  The requested operation is not allowed on a loopback target.
- **OVERWRITE_SNAPSHOT_BAD_NAME**

  The snapshot name does not exist.
- **OVERWRITE_SNAPSHOT_GROUP_DOES_NOT_BELONG_TO_GIVEN_GROUP**

  The snapshot group belongs to another consistency group.
- **POOL_SNAPSHOT_LIMIT_REACHED**

  There is not enough space to create a snapshot.
- **REMOTE_POOL_SNAPSHOT_LIMIT_REACHED**

  There is not enough space on the remote target for creating a snapshot.
- **REMOTE_MAX_VOLUMES_REACHED**

  The maximum number of volumes on the remote machine is already reached.
- **REMOTE_MAX_SNAPSHOTS_FOR_VOLUME_REACHED**

  The maximum allowed number of snapshots per volume is already reached on a remote machine whose version is not 10.2.4.
- **REMOTE_VOLUME_IS_MASTER**

  A volume on the remote machine is already defined as primary.
- **REMOTE_VOLUME_IS_SNAPSHOT**

  The secondary volume is a snapshot.

- **REMOTE_VOLUME_DATA_MIGRATION_UNSYNCHRONIZED**

  Data Migration to the remote volume has not completed.
- **REMOTE_SNAPSHOT_NAME_EXISTS**

  The remote snapshot name already exists.
- **REMOTE_SNAPSHOT_ILLEGAL_PRIORITY**

  Illegal snapshot priority (remote); must be an integer between 1 and 4.
- **REMOTE_SNAPSHOT_GROUP_NAME_EXISTS**

  The remote snapshot group name already exists.
- **REMOTE_SNAPSHOT_GROUP_ILLEGAL_PRIORITY**

  Illegal snapshot group priority (remote); must be an integer between 1 and 4.
- **REMOTE_SNAPSHOT_GROUP_BAD_PREFIX**

  The remote snapshot group name has a reserved prefix.
- **REMOTE_SNAPSHOT_BAD_PREFIX**

  The remote snapshot name has a reserved prefix.
- **REMOTE_CONS_GROUP_EMPTY**

  The operation is not allowed on an empty consistency group (remote).
- **REMOTE_CONS_GROUP_MISMATCH**

  The remote snapshot group does not match the consistency group volumes.
- **SNAPSHOT_HAS_ACTIVE_SYNC_JOB**

  The snapshot is currently the target of an active sync job.

  **Troubleshooting:** Please wait for the sync job to complete.
- **SNAPSHOT_ILLEGAL_PRIORITY**

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

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

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

  The snapshot group name already exists.
- **SNAPSHOT_GROUP_ILLEGAL_PRIORITY**

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

  The snapshot group name does not exist.
- **SNAPSHOT_GROUP_BAD_PREFIX**

  The snapshot group name has a reserved prefix.
- **SNAPSHOT_IS_PART_OF_SNAPSHOT_GROUP**

  The snapshot is part of a snapshot group.
- **SYNCHED_SNAPSHOTS_NOT_SUPPORTED_IN_TARGET**

  The mirror's target does not support the synchronized snapshot capability.
- **VOLUME_BAD_PREFIX**

  The volume name has a reserved prefix.

- **VOLUME_BELONGS_TO_MIRRORED_CONS_GROUP**

  The volume mirror is part of a consistency group mirror.
- **VOLUME_DATA_MIGRATION_UNSYNCHRONIZED**

  Data Migration to this volume has not completed.
- **VOLUME_EXISTS**

  The volume name already exists.
- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **VOLUME_NO_MIRROR**

  The local volume does not have remote mirroring definitions.
- **VOLUME_IS_NOT_CONSISTENT_SLAVE**

  The operation not allowed on an inconsistent secondary volume.
- **VOLUME_IS_SNAPSHOT**

  The operation is not permitted on snapshots.
- **VOLUME_IS_OLVM_PROXY**

  The volume is in an IBM Hyper-Scale Mobility Proxy phase.
- **OPERATION_DENIED_OBJECT_MANAGED**

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

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

  The target name does not match the consistency group.
- **TARGET_BAD_NAME**

  The target name does not exist.
- **VOLUME_TARGET_MISMATCH**

  The volume and target do not match.
- **REMOTE_MIRROR_IS_STANDBY**

  The remote mirror is marked as Standby.
- **VOLUME_HAS_MULTIPLE_MIRRORS**

  The volume has multiple mirrors. The operation is not allowed or a target must be specified.
- **CONS_GROUP_HAS_MULTIPLE_MIRRORS**

  The consistency group has multiple mirrors. The operation is not allowed, or a target must be specified.
- **MIRROR_IS_STANDBY**

  The mirror is marked as Standby.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support
- **MAX_SNAPSHOTS_PER_VOLUME_REACHED**

  The maximum allowed number of snapshots is already reached.

- **REMOTE_MAX_SNAPSHOTS_PER_VOLUME_REACHED**

  The maximum allowed number of snapshots is already reached on the remote ssytem.
- **TARGET_SNAPSHOT_GROUP_BAD_NAME**

  The target snapshot group name does not exist.
- **REMOTE_DOMAIN_MAX_VOLUMES_REACHED**

  The maximum number of volumes in the remote machine domain is already reached.
- **REMOTE_CONS_GROUP_BAD_NAME**

  The remote consistency group name does not exist.
- **SNAPSHOT_CAN_NOT_BE_CREATED_REMOTE_CONS_GROUP_IO_IS_NOT_PAUSED**

  The snapshot group will not be created since the remote consistency group is not in a stopped state.
- **SNAPSHOT_CAN_NOT_BE_CREATED_REMOTE_CONS_GROUP_DEFINITION_CHANGED**

  The snapshot group will not be created since the volumes in the remote consistency group have changed since the io_pause command was issued.
- **REMOTE_OVERWRITE_SNAPSHOT_GROUP_DOES_NOT_BELONG_TO_GIVEN_GROUP**

  The remote snapshot group belongs to another consistency group.
- **REMOTE_SNAPSHOT_IS_PART_OF_SNAPSHOT_GROUP**

  The subordinate snapshot is part of a snapshot group.
- **REMOTE_SNAPSHOT_IS_INTERNAL**

  Internal snapshot cannot be mapped, modified or deleted.
- **REMOTE_SNAPSHOT_HAS_ACTIVE_SYNC_JOB**

  The subordinate snapshot is currently the target of an active sync job.

  **Troubleshooting:** Please wait for the sync job to complete.
- **REMOTE_OVERWRITE_SNAPSHOT_IS_MASTER_VOL**

  The subordinate snapshot cannot be overwritten because it is a primary volume.
- **REMOTE_OVERWRITE_SNAPSHOT_BAD_NAME**

  The subordinate snapshot name does not exist.
- **REMOTE_SNAPSHOT_OVERWRITE_MISMATCH**

  The specified subordinate snapshot is not a snapshot of the specified volume.
- **REMOTE_MULTISITE_IS_NOT_SMASTER**

  The remote peer is not the SMaster.
- **MULTISITE_BAD_GLOBAL_ID**

  The Multi-site global ID does not exist.
- **MULTISITE_SMASTER_INVALID_CONFIGURATION**

  The Multi-site SMaster-Master relation configuration is invalid.

## Activating mirroring

Use the **mirror_activate** command to activate mirroring for a defined mirror coupling.

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

## Parameters

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

This command activates the coupling - either volumes or consistency groups - and switches it to the Active state.

Requirements for a successful command completion:

- The specified target must exist
- The specified target must be mirrored
- The specified target is a volume that does not belong to a consistency group, or is a consistency group
- The specified target is not a master
- The Standby state was explicitly set by issuing the **mirror_deactivate** command on the same peer

If the new activation state is the same as the existing state, nothing is done and a success code is returned.

The mirroring cannot be activated:

- If the time stamps of the last_replicated snapshots on the master and slave do not match.
- If the command is issued on a master that did not receive acknowledgment from the slave following the **cg_add_vol** or **cg_remove_vol** command (due to the command's timeout or to an unexpected failure), the command fails and the MIRROR_CONS_GROUP_MEMBERSHIP_MISMATCH code is returned. It means that the member lists of the mirror consistency group peers are not the same.
- If the command is issued on a master that did not receive acknowledgment from the slave following a **vol_resize** command (due to the command's timeout or to an unexpected failure), the command fails and the MIRROR_CONS_GROUP_MEMBERSHIP_MISMATCH code is returned. It means that the sizes of the mirror volume peers are not the same.

## Access control

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

## Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **VOLUME_NO_MIRROR**

The local volume does not have remote mirroring definitions.

- **CONS_GROUP_BAD_NAME**

  The consistency group name does not exist.

- **CONS_GROUP_NO_MIRROR**

  The local consistency group does not have remote mirroring definitions.

- **LOCAL_PEER_IS_NOT_MASTER**

  The local peer is not primary.

- **MIRROR_CONFIGURATION_ERROR**

  The mirror's local configuration does not match its remote configuration.

- **REMOTE_MAX_VOLUMES_REACHED**

  The maximum number of volumes on the remote machine is already reached.

- **SYNC_ALREADY_ACTIVE**

  Synchronization is already active.

- **VOLUME_BELONGS_TO_MIRRORED_CONS_GROUP**

  The volume mirror is part of a consistency group mirror.

- **DOMAIN_MAX_VOLUMES_REACHED**

  The domain exceeds the maximum allowed number of volumes.

- **MIRROR_CAN_NOT_BE_ACTIVATED**

  Mirroring cannot be activated.

- **MIRROR_CONS_GROUP_MEMBERSHIP_MISMATCH**

  The mirrored consistency group contains different volumes on the primary and secondary machines. This problem occurs whenever the cg_add_vol or cg_remove_vol commands were previously issued, and the primary machine did not receive an acknowledgment from the secondary machine until the command timed out, or due to any other unexpected failure.

- **MIRROR_SIZE_MISMATCH**

  The secondary and primary volume sizes are different.

- **MIRROR_RETRY_OPERATION**

  There is an operation in progress on this mirror.

  **Troubleshooting:** Retry the command in a few seconds.

- **VOLUME_HAS_MULTIPLE_MIRRORS**

  The volume has multiple mirrors. The operation is not allowed or a target must be specified.

- **CONS_GROUP_HAS_MULTIPLE_MIRRORS**

  The consistency group has multiple mirrors. The operation is not allowed, or a target must be specified.

- **TARGET_BAD_NAME**

  The target name does not exist.

- **VOLUME_TARGET_MISMATCH**

  The volume and target do not match.

- **CONS_GROUP_BAD_TARGET**

  The target name does not match the consistency group.

- **VOLUME_TOO_MANY_ACTIVE_MIRRORS**

  This command cannot be issued if more than one mirror is active on the volume.

- **REMOTE_MIRROR_IS_STANDBY**

The remote mirror is marked as Standby.

- **REMOTE_DOMAIN_MAX_VOLUMES_REACHED**

  The maximum number of volumes in the remote machine domain is already reached.
- **MIRROR_IS_STANDBY**

  The mirror is marked as Standby.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support
- **REMOTE_DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier of the remote system is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support
- **REMOTE_SYSTEM_OUT_OF_PHYSICAL_SPACE**

  The operation not allowed while the remote system is out of physical space.
- **CONS_GROUP_MEMBER_VOL_IS_MISSING_A_RELATION**

  A volume in the consistency group is missing the volume-level Multi-site/mirror/HyperSwap relation.
- **REMOTE_MULTISITE_IS_NOT_SMASTER**

  The remote peer is not the SMaster.
- **MULTISITE_SMASTER_INVALID_CONFIGURATION**

  The Multi-site SMaster-Master relation configuration is invalid.
- **MULTISITE_BAD_GLOBAL_ID**

  The Multi-site global ID does not exist.

## Changing the RPO for local or remote system

Use the **mirror_change_rpo** command to change a local or remote RPO for a mirror relation.

```
mirror_change_rpo  <vol=VolName | cg=cgName> [ target=TargetName ] [ rpo=rpo ] [ remote_rpo=rpo ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **vol** | Object name | Local volume name.<br>Must be specified if the command is applied to a volume. | N | N/A |
| **cg** | Object name | Consistency group name on the local system. | N | N/A |
| **target** | Object name | Target name of the mirror, mandatory if there are 2 mirrors defined on the volume. | N | [none] |
| **remote_rpo** | Integer | RPO on a remote system. | N | [Unchanged] |
| **rpo** | Integer | RPO on the local system | N | [Unchanged] |

- The command must be run on the master.

- The RPO must be greater than the interval.
- The link has to be up.

**Example:**

```
mirror_change_rpo vol=volname rpo=100
```

**Output:**

```
Command completed successfully.
```

## Access control

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

## Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **CONS_GROUP_NO_MIRROR**

  The local consistency group does not have remote mirroring definitions.
- **ASYNC_MIRROR_REMOTE_RPO_TOO_SHORT**

  The specified remote RPO is too short.
- **ASYNC_MIRROR_RPO_TOO_LONG**

  The specified RPO is too long.
- **TARGET_NOT_CONNECTED**

  There is currently no connection to the target system.
- **VOLUME_NO_MIRROR**

  The local volume does not have remote mirroring definitions.
- **CONS_GROUP_BAD_NAME**

  The consistency group name does not exist.
- **ASYNC_MIRROR_RPO_TOO_SHORT**

  The specified RPO is too short.
- **VOLUME_BELONGS_TO_MIRRORED_CONS_GROUP**

  The volume mirror is part of a consistency group mirror.
- **INTERVAL_SHOULD_BE_SHORTER_THAN_RPO**

  The schedule interval must be shorter than the RPO.
- **ASYNC_MIRROR_REMOTE_RPO_TOO_LONG**

  The specified remote RPO is too long.

- **LOCAL_IS_SLAVE**

  The local mirror peer is not primary.
- **SYNC_MIRROR_HAS_NO_RPO**

  The synchronous mirror does not have an RPO.
- **TARGET_BAD_NAME**

  The target name does not exist.
- **VOLUME_HAS_MULTIPLE_MIRRORS**

  The volume has multiple mirrors. The operation is not allowed or a target must be specified.
- **CONS_GROUP_HAS_MULTIPLE_MIRRORS**

  The consistency group has multiple mirrors. The operation is not allowed, or a target must be specified.
- **VOLUME_TARGET_MISMATCH**

  The volume and target do not match.
- **CONS_GROUP_BAD_TARGET**

  The target name does not match the consistency group.
- **MIRROR_IS_STANDBY**

  The mirror is marked as Standby.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support

## Changing the designation of mirroring peers

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

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

### Parameters

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

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

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

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

**Example:**

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

**Output:**

```
Command completed successfully.
```

## Access control

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

## Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **VOLUME_NO_MIRROR**

  The local volume does not have remote mirroring definitions.
- **CONS_GROUP_BAD_NAME**

  The consistency group name does not exist.
- **CONS_GROUP_NO_MIRROR**

  The local consistency group does not have remote mirroring definitions.
- **LOCAL_PEER_IS_NOT_MASTER**

  The local peer is not primary.
- **MIRROR_DESIGNATION_NOT_SUPPORTED_BY_TARGET**

  The mirror's target does not support mirror role designation.
- **MIRROR_IS_NON_OPERATIONAL**

  The mirror is non-operational.
- **VOLUME_BELONGS_TO_MIRRORED_CONS_GROUP**

  The volume mirror is part of a consistency group mirror.
- **TARGET_BAD_NAME**

  The target name does not exist.
- **VOLUME_HAS_MULTIPLE_MIRRORS**

  The volume has multiple mirrors. The operation is not allowed or a target must be specified.
- **CONS_GROUP_HAS_MULTIPLE_MIRRORS**

  The consistency group has multiple mirrors. The operation is not allowed, or a target must be specified.

- **VOLUME_TARGET_MISMATCH**

  The volume and target do not match.
- **CONS_GROUP_BAD_TARGET**

  The target name does not match the consistency group.
- **MIRROR_IS_STANDBY**

  The mirror is marked as Standby.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support

## Changing a mirroring schedule for remote slave peers

Use the **mirror_change_remote_schedule** command to change the replication schedule of a remote slave peer.

```
mirror_change_remote_schedule < vol=VolName | cg=cgName > [ target=TargetName ]
remote_schedule=Schedule
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **vol** | Object name | Local master volume name. | N | N/A |
| **cg** | Object name | Local master consistency group name. | N | N/A |
| **target** | Object name | Target mirror name. Mandatory if 2 mirrors are defined on the volume. | N | [none] |
| **remote_schedule** | Object name | A reference to a remote schedule that should be set for the remote slave peer, which corresponds with the master specified in the command. | Y | N/A |

This command changes the replication schedule of an asynchronous coupling in order to make it effective after the role of a specified remote slave peer is changed to master.

Prerequisites:

- The coupling must be ASYNC_INTERVAL.

Following the command execution:

- The system displays a warning
- If the command is approved, it is executed
- An event is generated
- New sync jobs are generated according to the updated schedule
- Existing sync jobs are not affected (that is, they run according to the previous schedule)

Requirements for a successful command completion:

- The specified target exists

- The specified target is mirrored
- The specified target is not a volume that belongs to a mirrored consistency group
- The specified target is of sync type ASYNC_INTERVAL
- The specified target is a master
- The link is up

## Access control

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

## Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **VOLUME_NO_MIRROR**

  The local volume does not have remote mirroring definitions.
- **CONS_GROUP_BAD_NAME**

  The consistency group name does not exist.
- **CONS_GROUP_NO_MIRROR**

  The local consistency group does not have remote mirroring definitions.
- **REMOTE_VOLUME_IS_MASTER**

  A volume on the remote machine is already defined as primary.
- **REMOTE_CONS_GROUP_IS_MASTER**

  The remote consistency group is defined as primary.
- **SCHEDULE_DOES_NOT_EXIST**

  The specified schedule does not exist.
- **SYNC_MIRROR_DOES_NOT_USE_SCHEDULE**

  A synchronous mirror definition does not require a schedule object.
- **TARGET_NOT_CONNECTED**

  There is currently no connection to the target system.
- **VOLUME_BELONGS_TO_MIRRORED_CONS_GROUP**

  The volume mirror is part of a consistency group mirror.
- **INTERVAL_SHOULD_BE_SHORTER_THAN_RPO**

  The schedule interval must be shorter than the RPO.
- **TARGET_BAD_NAME**

  The target name does not exist.
- **VOLUME_HAS_MULTIPLE_MIRRORS**

The volume has multiple mirrors. The operation is not allowed or a target must be specified.

- **CONS_GROUP_HAS_MULTIPLE_MIRRORS**

  The consistency group has multiple mirrors. The operation is not allowed, or a target must be specified.
- **VOLUME_TARGET_MISMATCH**

  The volume and target do not match.
- **CONS_GROUP_BAD_TARGET**

  The target name does not match the consistency group.
- **DOMAIN_HAS_NO_ACCESS_TO_SCHEDULE**

  The domain has no access to the schedule.
- **MIRROR_IS_STANDBY**

  The mirror is marked as Standby.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support

## Changing the role of a mirrored volume

Use the **mirror_change_role** command to change the role of a local mirroring peer from Master to Slave or from Slave to Master.

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

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **vol** | Object name | Local volume name.<br>Must be specified if the command is applied to a volume. | N | N/A |
| **cg** | Object name | Consistency group name.<br>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.<br>If not specified, the command swaps peer roles between Master and Slave. | N | none |

This command changes the role of the local peer from Master to Slave or from Slave to Master when the coupling is non-operational. It is assumed that the command will be issued on both peers of the coupling before the coupling becomes operational again, so that upon reconnection there still will be one Master and one Slave.

For a successful command completion:

- Do not issue the command on a peer whose status is Initializing.
- Do not issue the command in the Change Tracking state.
- Do not issue the command on a volume that belongs to a mirrored consistency group, otherwise the command will return an error and fail.

Changing the roles in synchronous mirroring:

When applied on a Master, the Master becomes a Slave, ceases serving host requests, and is set to accept replication from the other peer as a Slave.

When applied on a Slave, the Slave becomes a Master, starts accepting requests from hosts, and upon explicit activation starts replicating to the other peer (the original Master).

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

Changing the roles in asynchronous mirroring:

When successfully applied on a Master, the Master is reverted to the image recorded on the last_replicated snapshot of the mirror, and ceases accepting host requests.

When applied on a Slave:

- A warning is displayed.
- An event is generated.
- The new Master ceases accepting replication requests from the previous Master, and reverts to the last_replicated snapshot.
- The new Master starts accepting host requests.
- The process completion is recorded in the log.
- Explicit activation of mirroring is required.

## Access control

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

## Warnings

- **SOME_DATA_WILL_BE_LOST_ARE_YOU_SURE**

  Are you sure you want the mirror/HyperSwap local peer to become secondary and lose the data that was not replicated?

- **ARE_YOU_SURE_YOU_WANT_TO_CHANGE_A_PEER_WITH_LCS_TO_MASTER**

  Are you sure you want the mirror/HyperSwap local peer to become primary? The local peer has a last-consistent snapshot.

## Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **VOLUME_NO_MIRROR**

  The local volume does not have remote mirroring definitions.
- **CONS_GROUP_BAD_NAME**

  The consistency group name does not exist.
- **CONS_GROUP_NO_MIRROR**

  The local consistency group does not have remote mirroring definitions.
- **MIRROR_IS_INITIAL**

  The operation is not permitted during the Initialization phase.
- **MIRROR_IS_ACTIVE**

  Remote mirroring is currently active.
- **VOLUME_HAS_DATA_MIGRATION**

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

  The volume mirror is part of a consistency group mirror.
- **MIRROR_RETRY_OPERATION**

  There is an operation in progress on this mirror.

  **Troubleshooting:** Retry the command in a few seconds.
- **MIRROR_HAS_NO_SYNCHED_SNAPSHOT**

  The mirror does not have a synchronized snapshot.
- **MASTER_CANNOT_BE_DEMOTED**

  The primary volume cannot be demoted to secondary. Peer status mismatch.
- **VOLUME_HAS_MULTIPLE_RELATIONS**

  The volume has multiple relations. The operation is not allowed or a target must be specified.
- **CONS_GROUP_HAS_MULTIPLE_RELATIONS**

  The consistency group has multiple relations. The operation is not allowed or a target must be specified.
- **TARGET_BAD_NAME**

  The target name does not exist.
- **VOLUME_TARGET_MISMATCH**

  The volume and target do not match.
- **CONS_GROUP_BAD_TARGET**

  The target name does not match the consistency group.
- **MIRROR_PART_OF_MULTISITE**

  The remote mirror is part of a Multi-site.
- **MIRROR_IS_STANDBY**

  The mirror is marked as Standby.
- **COMMAND_NOT_SUPPORTED_FOR_OLVM_VOLUMES**

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

- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support
- **SYSTEM_OUT_OF_PHYSICAL_SPACE**

  The operation not allowed while the system is out of physical space.

# Changing a mirroring schedule for local peers

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

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

## Parameters

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

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

Prerequisites:

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

The command fails under the following conditions:

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

Setting a scheduling reference:

- An event is generated
- New sync jobs will be generated according to updated schedule. A running sync job is unaffected.

## Access control

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

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

## Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **CONS_GROUP_BAD_NAME**

  The consistency group name does not exist.
- **VOLUME_NO_MIRROR**

  The local volume does not have remote mirroring definitions.
- **CONS_GROUP_NO_MIRROR**

  The local consistency group does not have remote mirroring definitions.
- **SCHEDULE_DOES_NOT_EXIST**

  The specified schedule does not exist.
- **SYNC_MIRROR_DOES_NOT_USE_SCHEDULE**

  A synchronous mirror definition does not require a schedule object.
- **VOLUME_BELONGS_TO_MIRRORED_CONS_GROUP**

  The volume mirror is part of a consistency group mirror.
- **INTERVAL_SHOULD_BE_SHORTER_THAN_RPO**

  The schedule interval must be shorter than the RPO.
- **ILLEGAL_INTERVAL**

  The specified interval value is not supported.
- **TARGET_BAD_NAME**

  The target name does not exist.
- **VOLUME_HAS_MULTIPLE_MIRRORS**

  The volume has multiple mirrors. The operation is not allowed or a target must be specified.
- **CONS_GROUP_HAS_MULTIPLE_MIRRORS**

  The consistency group has multiple mirrors. The operation is not allowed, or a target must be specified.
- **VOLUME_TARGET_MISMATCH**

  The volume and target do not match.
- **CONS_GROUP_BAD_TARGET**

  The target name does not match the consistency group.
- **MIRROR_IS_STANDBY**

  The mirror is marked as Standby.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support

# Creating a mirroring definition

Use the **mirror_create** command to create a remote mirroring coupling.

```
mirror_create < vol=VolName slave_vol=SlaveVolumeName [ create_slave=<yes|no>
[ remote_pool=RemotePoolName ] ]
 [ init_type=<online|offline> ]
 > | <cg=cgName slave_cg=SlaveCgName> [ type=<SYNC_BEST_EFFORT|ASYNC_INTERVAL> ]
 target=TargetName [ rpo=rpo [ remote_rpo=rpo ] schedule=Schedule remote_schedule=Schedule ]
 [ part_of_multisite=<yes|no> ]
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **vol** | Object name | Local volume to be mirrored (the master). | N | N/A |
| **slave_vol** | Object name | The name of the slave volume on the remote storage system. | N | N/A |
| **create_slave** | Boolean | Determines whether to create a new slave volume or to use an existing one. | N | no |
| **remote_pool** | Object name | The storage pool on the remote system. Relevant only if creating a slave. | N | N/A |
| **cg** | Object name | Local consistency group to be mirrored (the master). | N | N/A |
| **slave_cg** | Object name | The name of the slave consistency group on the remote storage system. | N | N/A |
| **type** | Enumeration | The name of the replication type | N | SYNC_BEST_ EFFORT |
| **target** | Object name | Remote target to contain the slave volume. | Y | N/A |
| **rpo** | Positive integer | A mirror recovery point objective value for the master. Ranges from 30 to 86400 seconds (that is, up to 24 hours)<br>Is applicable and mandatory for asynchronous mirroring only. | N | [None] |
| **remote_rpo** | Positive integer | Mirror recovery point objective value for a remote peer that becomes master<br>Is applicable and mandatory for asynchronous mirroring only. | N | [Master RPO] |

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **schedule** | Object name | A reference to a schedule object<br>Is applicable and mandatory for asynchronous mirroring only. | N | [None] |
| **remote_schedule** | Object name | A reference to a schedule object on the remote machine.<br>Is applicable and mandatory for asynchronous mirroring only. | N | [None] |
| **init_type** | Enumeration | Specifies the method requested to initialize the slave mirror. | N | [none] |
| **part_of_multisite** | Boolean | Marks the mirror as part of a Multi-site relation. | N | no |

Mirroring is the process of ensuring that both peers contain identical data at all times. This command defines a new mirroring coupling between a master and a slave peers.

The command supports the creation of an asynchronous mirroring coupling. Asynchronous mirroring is based on schedule-driven replication. The system also offers a predefined schedule object with a non-user-configurable interval of 20 seconds, named **min_interval**.

To create a mirroring coupling, an existing master peer must be specified together with a slave peer. Upon creation, the coupling is not active and the user needs to activate it explicitly in order to start the replication. This slave either already exists or is created by this command. Using an existing slave is allowed only if it is formatted. If the slave already exists, the command receives its name along with the remote system name. If it is created by this command, the input parameters specify the remote storage system name, the name of the slave that is created and the storage pool that will contain the newly created slave.

To add a second mirror (Multi-site) for an existing mirrored volume, use the **part_of_multisite** flag for the new mirror.

Mirroring is created in the standby state. The mirroring coupling must then be activated in order to start the initialization process, which copies the data from the master to the slave.

A storage system can have multiple mirroring definitions between pairs of peers on various remote systems. However, when the peers are consistency groups, all the volumes included in a specific consistency group must be mirrored between only one pair of storage systems. Therefore, when a volume peer on a storage system (for example: A) has a mirroring relation with a volume on a remote storage system (for example: B), any other volume in the same consistency group on storage system A can only be defined in a remote mirroring relation with a volume on storage system B. The same goes for volumes from storage system B to A. In addition, the mirrored consistency group has one sync job for all pertinent mirrored volumes within the consistency group.

The command fails if it finds conflicting mirroring snapshots (that were not removed during the deletion of a previous mirroring definition).

Initialization types:

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

To create a remote mirroring coupling as part of a Multi-site relation, apply the **part_of_multisite** parameter.

## Access control

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

## Warnings

- **VOLUME_SIZE_VERY_LARGE_ARE_YOU_SURE**

  The volume size is very large. It may not be possible to mirror this volume to older versions of the storage system. Are you sure?

## Return codes

- **ASYNC_MIRROR_MISSING_RPO**

  An asynchronous mirror definition must include the RPO.
- **ASYNC_MIRROR_REMOTE_RPO_TOO_LONG**

  The specified remote RPO is too long.
- **ASYNC_MIRROR_REMOTE_RPO_TOO_SHORT**

  The specified remote RPO is too short.
- **ASYNC_MIRROR_RPO_TOO_SHORT**

  The specified RPO is too short.
- **ASYNC_MIRROR_RPO_TOO_LONG**

  The specified RPO is too long.
- **ASYNC_NOT_SUPPORTED_IN_TARGET**

  The specified target does not support asynchronous mirroring.
- **BAD_REMOTE_VOLUME_NAME**

  The secondary volume name does not exist.
- **BAD_REMOTE_VOLUME_SIZE**

  The primary and secondary volumes contain a different number of blocks.
- **CONS_GROUP_BAD_NAME**

  The consistency group name does not exist.
- **CONS_GROUP_HAS_MIRROR**

  Mirroring is defined for this consistency group.
- **CONS_GROUP_MIRRORING_NOT_SUPPORTED_IN_TARGET**

  Consistency group mirroring is not supported by the target machine.
- **INTERVAL_SHOULD_BE_SHORTER_THAN_RPO**

  The schedule interval must be shorter than the RPO.

- **ILLEGAL_INTERVAL**

  The specified interval value is not supported.
- **MAX_MIRRORS_REACHED**

  The maximum number of mirrors is already reached.
- **MAX_SYNC_MIRRORS_REACHED**

  The maximum number of sync mirrors is already defined
- **MAX_ASYNC_MIRRORS_REACHED**

  The maximum number of async mirrors is already reached.
- **NOT_ENOUGH_SPACE_ON_REMOTE_MACHINE**

  Not enough free space to set the requested size of the secondary volume.
- **NO_ASYNC_IN_THIN_PROVISIONED_POOL**

  A thin-provisioned pool cannot contain volumes with asynchronous mirroring.
- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **VOLUME_IS_MASTER**

  This local volume is already defined as a primary volume.
- **VOLUME_IS_SLAVE**

  The volume is defined as a secondary volume.
- **REMOTE_VOLUME_EXISTS**

  The secondary volume with the indicated name already exists. The name cannot be reused.
- **REMOTE_MAX_VOLUMES_REACHED**

  The maximum number of volumes on the remote machine is already reached.
- **REMOTE_MAX_MIRRORS_REACHED**

  The maximum number of mirrors is already defined on remote machine
- **REMOTE_CONS_GROUP_NOT_EMPTY**

  The remote consistency group contains volumes.
- **VOLUME_BAD_PREFIX**

  The volume name has a reserved prefix.
- **REMOTE_POOL_DOES_NOT_EXIST**

  The pool does not exist on the remote machine.
- **REMOTE_POOL_NOT_SPECIFIED**

  Prior to creating a secondary volume, a pool must be defined on the remote machine.
- **REMOTE_TARGET_NOT_CONNECTED**

  There is currently no connection from the target system.
- **VOLUME_IS_SNAPSHOT**

  The operation is not permitted on snapshots.
- **REMOTE_VOLUME_IS_SNAPSHOT**

  The secondary volume is a snapshot.
- **TARGET_BAD_NAME**

  The target name does not exist.
- **TARGET_BAD_TYPE**

The target machine is not an XIV machine.

- **TARGET_NO_ACCESS**

  No access permissions to the secondary machine.
- **TARGET_NOT_CONNECTED**

  There is currently no connection to the target system.
- **REMOTE_VOLUME_LOCKED**

  The secondary volume is locked.
- **TIMEOUT**

  A remote operation was not completed in time.
- **VOLUME_HAS_MIRRORING_SNAPSHOTS**

  The volume has snapshots created by a previous mirroring process.
- **SLAVE_VOLUME_NOT_FORMATTED**

  The secondary volume is not formatted.
- **TARGET_DOES_NOT_ACCEPT_XIV_COMMANDS**

  The target system does not accept XIV management commands.
- **SYNC_MIRROR_HAS_NO_RPO**

  The synchronous mirror does not have an RPO.
- **REMOTE_CONS_GROUP_IS_MIRRORED**

  Mirroring is defined for this remote consistency group.
- **REMOTE_SCHEDULE_DOES_NOT_EXIST**

  The specified schedule does not exist on the remote machine.
- **SCHEDULE_DOES_NOT_EXIST**

  The specified schedule does not exist.
- **REMOTE_CONS_GROUP_BAD_NAME**

  The remote consistency group name does not exist.
- **REMOTE_VOLUME_IS_MASTER**

  A volume on the remote machine is already defined as primary.
- **REMOTE_VOLUME_IS_SLAVE**

  A volume on the remote machine is already defined as secondary.
- **REMOTE_MAX_MIRROR_CAPACITY_REACHED**

  The maximum capacity for mirrored volumes is already reached on the remote machine.
- **MIRROR_RETRY_OPERATION**

  There is an operation in progress on this mirror.

  **Troubleshooting:** Retry the command in a few seconds.
- **MIRRORING_INCOMPATIBLE_TARGET_VERSION**

  Mirroring is not supported between the system versions of the specified peers.
- **MIRROR_TYPE_INCOMPATIBLE_WITH_TARGET**

  A mirror of this type is not supported between the system versions of the specified peers.
- **NO_OFFLINE_INIT_TYPE_WITH_SLAVE_CREATION**

  A new volume will be created as secondary. Offline initialization is meaningless.
- **ASYNC_WITH_OFFLINE_INIT_NOT_SUPPORTED_IN_TARGET**

The specified target does not support asynchronous mirroring with offline initialization.

- **VOLUME_SIZE_ABOVE_LIMIT**

  The specified volume size is above the limit.

- **REMOTE_VOLUME_SIZE_ABOVE_LIMIT**

  The specified volume size is above the limit of the remote machine.

- **INVALID_SLICE_OFFSET**

  Slice offset is illegal.

- **VOLUME_IS_OLVM_PROXY**

  The volume is in an IBM Hyper-Scale Mobility Proxy phase.

- **REMOTE_VOLUME_IS_OLVM_PROXY**

  The remote volume is in an IBM Hyper-Scale Mobility Proxy phase.

- **ENCRYPTION_IN_PROGRESS**

  The system is in the process of changing the encryption activation state.

- **MIRROR_OF_SAME_TYPE_EXISTS_ON_VOLUME**

  A mirror of this type is already defined on this volume.

- **MIRROR_EXISTS_ON_TARGET**

  The volume already has a mirror on this target.

- **REMOTE_VOLUME_IS_MIRROR_MASTER**

  The volume is primary in a mirror relationship, and cannot be secondary!

- **MULTISITE_MAX_NUM_OF_MIRRORS_REACHED**

  Failed to create the relation, the maximum allowed number of relations is already exceeded.

- **REMOTE_VOLUME_TWO_SYNC_MIRRORS_NOT_ALLOWED**

  Two synchronous mirrors were detected on the remote volume. This is not allowed.

- **REMOTE_VOLUME_MIRROR_LOOP_DETECTED**

  A mirror loop was detected on the remote volume. This means that there is a mirror on the remote system, whose target is this system. Therefore, you cannot create a mirror with this target on this system.

- **DOMAIN_MAX_MIRRORS_REACHED**

  The domain exceeds the maximum allowed number of mirrors.

- **REMOTE_DOMAIN_MAX_VOLUMES_REACHED**

  The maximum number of volumes in the remote machine domain is already reached.

- **REMOTE_DOMAIN_HAS_NO_ACCESS_TO_TARGET**

  The secondary machine domain has no access to the target.

- **REMOTE_DOMAIN_HAS_NO_ACCESS_TO_SCHEDULE**

  The secondary machine domain has no access to the schedule.

- **DOMAIN_HAS_NO_ACCESS_TO_TARGET**

  The domain has no access to the target.

- **REMOTE_DOMAIN_MAX_MIRRORS_REACHED**

  The maximum number of mirrors is already reached in the remote machine domain.

- **DOMAIN_MAX_VOLUMES_REACHED**

  The domain exceeds the maximum allowed number of volumes.

- **REMOTE_VOLUME_HAS_DATA_MIGRATION**

  Data migration is already defined for the secondary volume.
- **REMOTE_VOLUME_MASTER_ASYNC_MIRROR_DETECTED**

  An asynchronous primary mirror was detected on the remote volume. The operation not allowed.
- **MAX_MULTISITES_REACHED**

  The number of Multi-site objects exceeded the limit.
- **MULTISITE_INCOMPATIBLE_TARGET_VERSION**

  Multi-site is not supported between the system versions of the specified peers.
- **REMOTE_VOLUME_HAS_MIRRORING_SNAPSHOTS**

  The remote volume has snapshots created by a previous mirroring process.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support
- **SYSTEM_OUT_OF_PHYSICAL_SPACE**

  The operation not allowed while the system is out of physical space.
- **REMOTE_DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier of the remote system is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support
- **REMOTE_SYSTEM_OUT_OF_PHYSICAL_SPACE**

  The operation not allowed while the remote system is out of physical space.
- **SLAVE_VOLUME_NOT_SAME_TYPE**

  The primary and secondary volumes are not of the same type. Either one of them is compressed and the other is not, or they use different compression technologies.
- **CONS_GROUP_IS_SECONDARY**

  The consistency group is already defined as a secondary consistency group.
- **CONS_GROUP_IS_PRIMARY**

  The consistency group is already defined as a primary consistency group.
- **RELATION_OF_SAME_TYPE_EXISTS_ON_CONS_GROUP**

  A relation of this type is already defined on this consistency group.
- **REMOTE_VOLUME_BAD_POOL**

  The remote volume and remote consistency group belong to different storage pools.
- **REMOTE_VOLUME_BELONGS_TO_CONS_GROUP**

  The remote volume belongs to a consistency group.
- **CONS_GROUP_MIRROR_ROLE_MISMATCH**

  All volumes in a mirrored consistency group must have the same mirroring role.
- **CONS_GROUP_MIRROR_SCHEDULE_MISMATCH**

  All volumes in a mirrored consistency group must have the same mirroring schedule.
- **CONS_GROUP_RELATION_STATE_IS_NOT_INITIALIZING**

  This operation is only allowed when the matching consistency group relation state is 'Initializing'.
- **VOLUME_RELATION_IN_CG_CAN_NOT_BE_CREATED**

A second relation must first be added to the consistency group before it can be added to a member volume.

- **CONS_GROUP_RELATION_PART_OF_MULTISITE_MISMATCH**

  All volumes in a mirrored or HyperSwap consistency group must have the same part_of_multisite setting.
- **CONS_GROUP_NOT_EMPTY**

  This operation is only allowed on an empty consistency group.
- **CONS_GROUP_MIRROR_TARGET_MISMATCH**

  All volumes in a mirrored consistency group must have the same mirroring target.
- **CONS_GROUP_RELATION_INIT_TYPE_MISMATCH**

  All volumes in a mirrored or HyperSwap consistency group must have the same init type setting.
- **REMOTE_MULTISITE_CONS_GROUP_MEMBER_VOL_HAS_TOO_MANY_RELATIONS**

  A remote volume belonging to the consistency group has too many relations.
- **MULTISITE_CANNOT_CONTAIN_SYNC_MIRROR**

  A Multi-site relation cannot include a synchronous mirroring relation.
- **REMOTE_VOLUME_HAS_A_PART_OF_MULTISITE_RELATION_BUT_NO_MULTISITE**

  The target volume already has a 'part of Multi-site' relation. Creating a mirror relation on a volume with an existing 'part of Multi-site' relation is not allowed, unless the volume has a configured Multi-site relation.

  **Troubleshooting:** Define a Multi-site relation on this volume and retry the command.
- **MIRROR_MULTISITE_MISMATCH**

  The remote and local volumes are associated with different Multi-site relations.

# Deactivating mirroring

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

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

## Parameters

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

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

The command cannot be issued on a slave.

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

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

Deactivating a consistency group affects all of its volumes.

The command fails under the following conditions:

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

## Access control

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

## Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **VOLUME_NO_MIRROR**

  The local volume does not have remote mirroring definitions.
- **CONS_GROUP_BAD_NAME**

  The consistency group name does not exist.
- **CONS_GROUP_NO_MIRROR**

  The local consistency group does not have remote mirroring definitions.
- **LOCAL_PEER_IS_NOT_MASTER**

  The local peer is not primary.
- **SYNC_ALREADY_INACTIVE**

  Synchronization is already inactive.
- **VOLUME_BELONGS_TO_MIRRORED_CONS_GROUP**

  The volume mirror is part of a consistency group mirror.
- **MIRROR_RETRY_OPERATION**

  There is an operation in progress on this mirror.

  **Troubleshooting:** Retry the command in a few seconds.

- **TARGET_BAD_NAME**

  The target name does not exist.
- **VOLUME_HAS_MULTIPLE_MIRRORS**

  The volume has multiple mirrors. The operation is not allowed or a target must be specified.
- **CONS_GROUP_HAS_MULTIPLE_MIRRORS**

  The consistency group has multiple mirrors. The operation is not allowed, or a target must be specified.
- **VOLUME_TARGET_MISMATCH**

  The volume and target do not match.
- **CONS_GROUP_BAD_TARGET**

  The target name does not match the consistency group.
- **REMOTE_MIRROR_IS_STANDBY**

  The remote mirror is marked as Standby.
- **MIRROR_IS_STANDBY**

  The mirror is marked as Standby.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support

# Deleting a remote mirroring definition

Use the **mirror_delete** command to delete a remote mirroring coupling definition.

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

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **vol** | Object name | Local master volume name. | N | N/A |
| **cg** | Object name | Local master consistency group name. | N | N/A |
| **target** | Object name | Target mirror name. Mandatory if 2 mirrors are defined on the volume. | N | [none] |
| **force_on_slave** | Boolean | Forces the deletion of the remote mirroring coupling definition even of a slave. Deleting a remote mirroring definition can be forced on the slave peer only when it is in the initialization phase. | N | no |

When a coupling is initially created or after it is deactivated, it is in *standby* mode. Only a standby coupling can be deleted. The command can only be issued on the master.

After the remote mirroring is deleted, both peers are configured as *none*, meaning that they are no longer configured as either master or slave.

Only the remote mirroring coupling definition is deleted. Neither the volumes themselves, nor their snapshots are deleted.

The local object specified in the **vol** parameter, must be a master.

To delete a remote mirroring coupling, the communication must be established. If there is no communication, mirroring is only deleted on the master, and a configuration error appears on the slave once the communication resumes.

Command outcome:

- An event is generated
- Overall coupling statistics are captured
- The outstanding pertinent sync jobs are deleted
- The process completion is recorded in the log

Deleting the mirroring definition when the link is down:

- When the link is down, this command only deletes the mirroring definition on the master.
- To delete the mirroring definition from the slave:
  - Run the **mirror_change_role** command to turn the slave into the master
  - Run **mirror_delete**

The **force_on_slave** parameter:

- The parameter **force_on_slave** can be issued only if mirroring is in the initialization phase. In any other mode, the role can be changed to master and the peer mirror can be deleted.

## Access control

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

## Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_DELETE_CG_MIRRORING**

  Are you sure you want to delete the mirroring relations of the consistency group and of all volumes in the consistency group?

## Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **VOLUME_NO_MIRROR**

  The local volume does not have remote mirroring definitions.
- **CONS_GROUP_BAD_NAME**

  The consistency group name does not exist.
- **CONS_GROUP_NO_MIRROR**

The local consistency group does not have remote mirroring definitions.

- **LOCAL_PEER_IS_NOT_MASTER**

  The local peer is not primary.

- **MIRROR_IS_ACTIVE**

  Remote mirroring is currently active.

- **FORCE_DELETE_NOT_ALLOWED_ON_MASTER**

  Deletion needs to be forced on secondary mirrors only.

- **VOLUME_BELONGS_TO_MIRRORED_CONS_GROUP**

  The volume mirror is part of a consistency group mirror.

- **MIRROR_RETRY_OPERATION**

  There is an operation in progress on this mirror.

  **Troubleshooting:** Retry the command in a few seconds.

- **MIRROR_IS_NOT_INITIALIZING**

  The operation is permitted only during the Initialization phase.

- **TARGET_BAD_NAME**

  The target name does not exist.

- **VOLUME_HAS_MULTIPLE_MIRRORS**

  The volume has multiple mirrors. The operation is not allowed or a target must be specified.

- **CONS_GROUP_HAS_MULTIPLE_MIRRORS**

  The consistency group has multiple mirrors. The operation is not allowed, or a target must be specified.

- **VOLUME_TARGET_MISMATCH**

  The volume and target do not match.

- **MIRROR_ASSOCIATED_WITH_MULTISITE**

  This mirror is associated with a defined Multi-site, the operation is not allowed.

- **CONS_GROUP_BAD_TARGET**

  The target name does not match the consistency group.

- **REMOTE_MIRROR_IS_STANDBY**

  The remote mirror is marked as Standby.

- **MIRROR_IS_STANDBY**

  The mirror is marked as Standby.

- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support

- **MIRROR_CONS_GROUP_MEMBERSHIP_MISMATCH**

  The mirrored consistency group contains different volumes on the primary and secondary machines. This problem occurs whenever the cg_add_vol or cg_remove_vol commands were previously issued, and the primary machine did not receive an acknowledgment from the secondary machine until the command timed out, or due to any other unexpected failure.

# Viewing the mirroring status

Use the **mirror_list** command to list the status and configuration of mirroring couplings.

```
mirror_list [ < [ vol=VolName ] [ target=TargetName ] > | cg=cgName | < [ scope=<cg|volume> ]
 [ sync_type=<sync_best_effort|async_interval> ] > ] [ domain=DomainName ]
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **sync_type** | Enumeration | List type. The avaiable options are: SYNC_BEST_EFFORT, ASYNC_INTERVAL, or All (if no value is specified) | N | All (if no value is specified) |
| **scope** | Enumeration | List type: all mirrors, all volumes, all CGs | N | All (if no value is specified) |
| **vol** | Object name | Local volume name. | N | [none] |
| **cg** | Object name | Local consistency group name. | N | [none] |
| **target** | Object name | Remote target name. | N | [none] |
| **domain** | Object name | The domain name. | N | All Domains |

This command shows current configuration and status for the remote mirroring of volumes or consistency groups. Size/part/time to synchronize are unknown if this is the slave and connection is broken.

The following default parameters are shown:

- **Name**
- **Mirror Type** (sync_best_effort, async_interval)
- **Mirror Object**: CG or Volume
- **Role**: Master or Slave
- **Remote System**: target name
- **Remote Peer**: volume name
- **Active**: Yes or No
- **Status**: Initializing, Synchronized, Unsynchronized, Consistent, Inconsistent, RPO OK, RPO Lagging, or Change Tracking
- **Link Up**: Yes or No

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

- **Designation**: Primary or Secondary
- **Estimated Sync Time**: Estimated time to synchronization in seconds
- **Size To Synchronize**: Size to synchronize in MiB
- **Operational**: Yes or No
- **Sync Progress** (in %)
- **Mirror Error**: Explains why mirroring is deactivated: No_Error, Configuration_Error, Secondary_Pool_Exhausted, Master_Pool_Exhausted, or No_Thin_Provisioning_Resources
- **Schedule Name**
- **Last replicated Snapshot Time**: the value is presented in yyyy-mm-dd hh:mm:ss format
- **Specified RPO**: the value is presented in h:mm:ss format

The following deactivation reasons can be read from the output list (available only in XML output format):

- No_Error
- Configuration_Error
- Secondary_Pool_Exhausted
- Master_Pool_Exhausted
- Remote_And_Local_Volume_Size_Mismatch
- Cons_Group_Membership_Mismatch
- Possible_Remote_And_Local_Volume_Size_Mismatch
- Possible_Cons_Group_Membership_Mismatch
- No_Thin_Provisioning_Resources
- Peer_Status_Mismatch
- Temporarily_Deactivated_For_Upgrade
- Mirror_is_Standby
- Out_Of_Physical_Space
- Data_Reduction_Offline
- Secondary_Out_Of_Physical_Space
- Secondary_Data_Reduction_Offline
- HA_Slave_Failover
- Data_Service_failure
- Remote_Volume_Down
- SCSI_Reservation_Error

| Field ID | Field output | Description | Default position |
|---|---|---|---|
| **local_peer_name** | Name | N/A | 1 |
| **mirror_object** | Mirror Object | N/A | 3 |
| **designation** | Designation | N/A | N/A |
| **current_role** | Role | N/A | 4 |
| **target_name** | Remote System | N/A | 5 |
| **remote_peer_name** | Remote Peer | N/A | 6 |
| **active** | Active | N/A | 7 |
| **sync_state** | Status | N/A | 9 |
| **connected** | Link Up | N/A | 10 |
| **estimated_sync_time** | Est. remaining time (sec) | N/A | N/A |
| **size_to_synchronize** | Size To Sync (MiB) | N/A | N/A |
| **operational** | Operational | N/A | N/A |
| **sync_progress** | Sync Progress (%) | N/A | N/A |
| **mirror_error** | Mirror Error | No Error, Secondary pool exhausted, Configuration error or No thin provisioning resources | N/A |
| **sync_type** | Mirror Type | N/A | 2 |
| **schedule_name** | Schedule Name | N/A | N/A |
| **last_replicated_snapshot_time** | Last Replicated | N/A | N/A |

| Field ID | Field output | Description | Default position |
|---|---|---|---|
| **last_replicated_snapshot_exists** | Has Last Replicated Snapshot | N/A | N/A |
| **specified_rpo** | RPO | N/A | N/A |
| **remote_rpo** | Remote RPO | N/A | N/A |
| **crash_consistent** | Crash Consistency | N/A | N/A |
| **validate** | Validation | N/A | N/A |
| **is_standby** | Standby | N/A | 8 |
| **arch** | Remote Arch | N/A | N/A |

**Output:**

```
<command id="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 |

| User Category | Permission |
|---|---|
| Technicians | Disallowed |

## Obtaining statistics on past sync jobs

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

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

### Parameters

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

The command output includes:

- Date and time created
- Date and time started to run
- Date and time finished
- Job size (MiB)

Either a volume or consistency group must be specified.

| Field ID | Field output | Default position |
|---|---|---|
| **created_at** | Created | 1 |
| **started_at** | Started | 2 |
| **finished_at** | Finished | 3 |
| **job_size** | Job Size (MiB) | 4 |
| **duration** | Job Duration (Sec) | 5 |
| **avg_sync_rate** | Average Sync Rate (MB/sec) | 6 |

**Example:**

```
mirror_statistics_get vol=VolName
```

**Output:**

```
<job id="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**

  The volume name does not exist.
- **CONS_GROUP_NO_MIRROR**

  The local consistency group does not have remote mirroring definitions.
- **MIRROR_HAS_NO_STATISTICS**

  Job statistics were not collected for this mirror.
- **LOCAL_IS_SLAVE**

  The local mirror peer is not primary.
- **VOLUME_BELONGS_TO_MIRRORED_CONS_GROUP**

  The volume mirror is part of a consistency group mirror.
- **VOLUME_NO_MIRROR**

  The local volume does not have remote mirroring definitions.
- **CONS_GROUP_BAD_NAME**

  The consistency group name does not exist.
- **SYNC_MIRROR_HAS_NO_STATISTICS**

  Job statistics do not exist for the synchronous mirror.
- **TARGET_BAD_NAME**

  The target name does not exist.
- **VOLUME_HAS_MULTIPLE_MIRRORS**

  The volume has multiple mirrors. The operation is not allowed or a target must be specified.
- **CONS_GROUP_HAS_MULTIPLE_MIRRORS**

  The consistency group has multiple mirrors. The operation is not allowed, or a target must be specified.
- **VOLUME_TARGET_MISMATCH**

  The volume and target do not match.
- **CONS_GROUP_BAD_TARGET**

  The target name does not match the consistency group.
- **MIRROR_IS_STANDBY**

  The mirror is marked as Standby.

# Switching roles between master and slave

Use the **mirror_switch_roles** command to switch roles between master and slave volumes.

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

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **vol** | Object name | Local volume name. | N | N/A |
| **cg** | Object name | Local consistency group name. | N | N/A |
| **target** | Object name | N/A | N | [none] |

The command can only be issued if coupling is operational and only on the master. For synchronous mirroring it can only be issued when the coupling is synchronized; for asynchronous mirroring it can only be issued if there are no outstanding sync jobs and the volume and its last replicated snapshot are identical.

Following the execution of the command:

- The volume that was previously the master becomes the slave
- The volume that was previously the slave becomes the master

Before this command switches roles, the system stops accepting new writes to the local volume. With synchronous mirrors the system performs all pending writes, and only after all pending writes have been committed, the roles are switched.

After the command is executed, the mirror remains active.

## Access control

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

## Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_SWITCH_ROLES**

  Are you sure you want to switch the roles in this relation?

## Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **VOLUME_NO_MIRROR**

  The local volume does not have remote mirroring definitions.
- **CONS_GROUP_BAD_NAME**

The consistency group name does not exist.

- **CONS_GROUP_NO_MIRROR**

  The local consistency group does not have remote mirroring definitions.
- **LOCAL_PEER_IS_NOT_MASTER**

  The local peer is not primary.
- **MIRROR_IS_NON_OPERATIONAL**

  The mirror is non-operational.
- **MIRROR_IS_NOT_SYNCHRONIZED**

  The mirror is not synchronized.
- **VOLUME_HAS_DATA_MIGRATION**

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

  There is currently no connection from the target system.
- **VOLUME_BELONGS_TO_MIRRORED_CONS_GROUP**

  The volume mirror is part of a consistency group mirror.
- **MIRROR_HAS_SYNC_JOB**

  The operation is not permitted on a mirror with active sync jobs.
- **MIRROR_RETRY_OPERATION**

  There is an operation in progress on this mirror.

  **Troubleshooting:** Retry the command in a few seconds.
- **MIRROR_MASTER_DIFFERS_FROM_SLAVE**

  The mirror's primary volume was written to after the last replicated snapshot was taken.
- **REMOTE_MIRROR_IS_NOT_ACTIVE**

  Remote mirroring is currently inactive.
- **TARGET_BAD_NAME**

  The target name does not exist.
- **VOLUME_HAS_MULTIPLE_RELATIONS**

  The volume has multiple relations. The operation is not allowed or a target must be specified.
- **CONS_GROUP_HAS_MULTIPLE_RELATIONS**

  The consistency group has multiple relations. The operation is not allowed or a target must be specified.
- **VOLUME_TARGET_MISMATCH**

  The volume and target do not match.
- **CONS_GROUP_BAD_TARGET**

  The target name does not match the consistency group.
- **REMOTE_MIRROR_IS_STANDBY**

  The remote mirror is marked as Standby.
- **MIRROR_IS_STANDBY**

  The mirror is marked as Standby.
- **COMMAND_NOT_SUPPORTED_FOR_OLVM_VOLUMES**

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

The data reduced tier is offline, the operation is not allowed.

**Troubleshooting:** Contact IBM Support

- **SYSTEM_OUT_OF_PHYSICAL_SPACE**

  The operation not allowed while the system is out of physical space.

# Retrieving RPO thresholds

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

```
rpo_thresholds_get
```

| Field ID | Field output | Default position |
|---|---|---|
| **increase_percentage** | Increase Percentage | 1 |
| **increase_absolute** | Increase Absolute | 2 |

**Example:**

```
rpo_thresholds_get
```

**Output:**

```
Increase Percentage   Increase Absolute
--------------------- -------------------
100                   3600
```

## Access control

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

# Setting an RPO threshold

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

```
rpo_thresholds_set  [ increase_percentage=percentage ] [ increase_absolute=absolute ]
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **increase_percentage** | Integer | The threshold for RPO increase (in per cent), beyond which an event should be created. | N | none |

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **increase_absolute** | Integer | The threshold for RPO increase, beyond which an event should be created. | N | none |

**Example:**

```
rpo_thresholds_set increase_percentage=percentage
```

**Output:**

```
Command completed successfully.
```

### Access control

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

### Return codes

- **INVALID_RPO_THRESHOLD_PERCENTAGE**

  The values should be in the [1,10000] range.
- **INVALID_RPO_THRESHOLD_ABSOLUTE**

  The values should be in the [1,1000000] range.

## Changing the interval of a schedule

Use the **schedule_change** command to change the interval of a schedule.

```
schedule_change schedule=Schedule interval=IntervalSize [ domain=DomainList ]
```

### Parameters

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

This command updates the schedule definition. Such definition can be referenced to when specifying asynchronous mirroring couplings.

Limitation:

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

Outcome:

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

**Example:**

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

**Output:**

```
Command completed successfully.
```

## Access control

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

## Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_UPDATE_THE_SCHEDULE**

  Are you sure you want to update this schedule? This change will affect all the mirrors using that schedule.

## Return codes

- **SCHEDULE_DOES_NOT_EXIST**

  The specified schedule does not exist.
- **BAD_SCHEDULE_TIME_FORMAT**

  Time format for a schedule is HH:MM[:SS].
- **ILLEGAL_INTERVAL**

  The specified interval value is not supported.
- **SCHEDULE_CAN_NOT_BE_UPDATED**

  The specified schedule cannot be updated.

- **INTERVAL_SCHEDULE_REQUIRES_ONLY_ONE_INTERVAL**

  Only one interval can be defined in an interval schedule.
- **SCHEDULE_EXCLUDE_TIMES_NOT_REQUIRED**

  An exclusion period can be defined only if exclude_time is set.
- **ZERO_LENGTH_EXCLUSION_PERIOD**

  An exclusion period's start time must be different from its end time.
- **DOMAIN_SCHEDULE_IN_USE**

  The schedule is in use, and therefore cannot be moved to another domain.
- **DOMAIN_DOESNT_EXIST**

  The domain does not exist.
- **ILLEGAL_TIME_SLOT_SPACE**

  The defined value must be larger than the minimum time slot, but smaller than the interval.

# Creating a schedule object

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

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

## Parameters

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

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

Limitations:

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

The **type** parameter:

Prior to the introduction of this parameter, each asynchronous mirror could be configured with an automatic schedule, whose interval specified how often a replication point and the corresponding

replication process (sync job) should be automatically created. It was also possible to instruct the system to create a manual replication point and a corresponding sync job for a mirror using the dedicated CLI command **mirror_create_snapshot**. Finally, a single predefined schedule named *Never* with no interval settings was provided for mirrors that only required manual sync job creation.

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

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

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

Once set, the schedule type cannot be changed.

**Example:**

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

**Output:**

```
Command completed successfully.
```

## Access control

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

## Return codes

- **SCHEDULE_EXISTS**

  The schedule name exists.
- **BAD_SCHEDULE_TIME_FORMAT**

  Time format for a schedule is HH:MM[:SS].
- **MAX_SYNC_SCHEDULES_REACHED**

  The maximum number of schedule objects has been reached.
- **ILLEGAL_INTERVAL**

  The specified interval value is not supported.
- **INTERVAL_SCHEDULE_REQUIRES_ONLY_ONE_INTERVAL**

  Only one interval can be defined in an interval schedule.
- **ZERO_LENGTH_EXCLUSION_PERIOD**

  An exclusion period's start time must be different from its end time.
- **SCHEDULE_EXCLUDE_TIMES_NOT_REQUIRED**

An exclusion period can be defined only if exclude_time is set.

- **ONLY_INTERVAL_SCHEDULE_MAY_HAVE_EXCLUSIONS**

  An exclusion period may defined only for an interval schedule.

- **DOMAIN_DOESNT_EXIST**

  The domain does not exist.

- **ILLEGAL_TIME_SLOT_SPACE**

  The defined value must be larger than the minimum time slot, but smaller than the interval.

# Triggering a schedule

Use the **schedule_create_tick** command to trigger a schedule-equivalent event for the couplings with the specified schedule.

```
schedule_create_tick schedule=Schedule
```

## Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **schedule** | Object name | The name of an asynchronously mirrored schedule. | Y |

This command triggers a schedule-equivalent, interval-arrived event for couplings with the specified schedule.

- The command triggers a new sync job for asynchronous mirror definitions that are configured with the manual schedule specified by the command. The command triggers a simultaneous event for all mirrors with the specified schedule (and only whenever the schedule is of a non-interval type) which is equivalent to the 'new-interval-arrived' event triggered automatically by the system for a mirror (with a schedule of type interval).
- The command is different from **mirror_create_snapshot** whereas it is applied to mirrors that do not have an interval-based schedule. Thus, even though an event is triggered immediately (as with mirror_create snapshot), no sync job is created for a pertinent mirror with the specified schedule (in case such a mirror has an outstanding sync job, as one might expect for mirrors with an interval-based schedule, if a new interval arrives during an outstanding job).
- The event is triggered for all pertinent couplings at the same time.
- A warning is displayed, requiring a user confirmation.

## Access control

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

## Return codes

- **SCHEDULE_DOES_NOT_EXIST**

The specified schedule does not exist.

- **SCHEDULE_IS_NOT_MANUAL**

  The specified schedule does not allow an external trigger.

## Deleting a schedule object

Use the **schedule_delete** command to delete a schedule for replication.

```
schedule_delete schedule=Schedule
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **schedule** | Object name | The name of the schedule to be deleted. | Y |

This command deletes a schedule definition.

The command can be issued successfully only if the schedule specified is not referenced by a mirror coupling, or if it is not a pre-defined schedule (**min_interval**).

Outcome:

- The command will delete the specified schedule.

**Example:**

```
schedule_delete schedule=hourly
```

**Output:**

```
Command completed successfully.
```

### Access control

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

### Return codes

- **SCHEDULE_IS_ASSIGNED**

  The specified schedule is currently assigned to a mirror.
- **SCHEDULE_CAN_NOT_BE_DELETED**

  The specified schedule cannot be deleted.
- **SCHEDULE_DOES_NOT_EXIST**

  The specified schedule does not exist.

# Listing a schedule object

Use the **schedule_list** command to list the schedule properties for the specified coupling.

```
schedule_list [ schedule=Schedule ] [ domain=DomainName ]
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **schedule** | Object name | The name of the schedule. | N | All |
| **domain** | Object name | The domain name. | N | All Domains |

The following default parameters are listed:

- Name
- Interval

The following optional parameters can be listed:

- Predefined (is the schedule a predefined object)
- Last Tick (last timestamp the schedule was fired)

| Field ID | Field output | Default position |
|---|---|---|
| **name** | Name | 1 |
| **interval** | Interval | 2 |
| **predefined** | Predefined | N/A |

**Example:**

```
schedule_list
```

**Output:**

```
Name           Interval
never
min_interval   00:00:20
ASYNC_None_3   00:02:00
```

## Access control

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

# Renaming a schedule

Use the **schedule_rename** command to rename a schedule object.

```
schedule_rename schedule=Schedule new_name=Name
```

## Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **schedule** | Object name | The current name of the schedule. | Y |
| **new_name** | Object name | The new name for the schedule. | Y |

It is not possible to rename a predefined schedule.

## Access control

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

## Return codes

- **SCHEDULE_DOES_NOT_EXIST**

  The specified schedule does not exist.
- **SCHEDULE_NAME_EXISTS**

  The new schedule name already exists.
- **SCHEDULE_CAN_NOT_BE_UPDATED**

  The specified schedule cannot be updated.

# Viewing sync job status

Use the **sync_job_list** command to list the statuses of queued and running sync jobs for asynchronous couplings

```
sync_job_list [ vol=VolName | cg=cgName ] [ domain=DomainName ]
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **vol** | Object name | Local volume name. | N | [none] |
| **cg** | Object name | Local consistency group name. | N | [none] |
| **domain** | Object name | The domain name. | N | All Domains |

The following parameters are displayed:

- Mirroring coupling (volume/consistency group)
- Job state (provides state of the sync job transmission process, running from the primary to secondary peer):
  - new: sync job is initiating
  - ready: sync job is ready to transfer from primary to secondary peer
  - active: sync job is transferring newly written date to secondary peer from primary peer
  - done: sync job is complete - appears only for volumes whose mirror sync jobs have completed but under a consistency group mirror sync job that is still active
- Job type: scheduled, ad hoc
- Schedule (name of the referenced schedule object)
- Interval length (if applicable)
- Job size
- Job progress
- Date created
- Time created
- Date started to run
- Time started to run

| Field ID | Field output | Default position |
|---|---|---|
| **job_object** | Job Object | 1 |
| **mirror_peer** | Local Peer | 2 |
| **source_snap** | Source | 3 |
| **target_snap** | Target | 4 |
| **job_state** | State | 5 |
| **part_of_cg_job** | Part of CG | 6 |
| **job_type** | Job Type | 7 |
| **created_at** | Created | N/A |
| **started_at** | Started | N/A |

## Access control

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

# 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:** The input and output syntax of CLI commands uses the legacy terminology of "Master", "SMaster", and "Slave" volumes, which in any documentation except the CLI reference, are referred to as "Primary", "Secondary", and "Tertiary". This inconsistency is a necessary compromise, required to avoid changes to older CLI commands that are in customer use, and also to keep the CLI terminology consistent across the board. The new terminology helps emphasize the commonality between the more recent functions of Multi-site HA/DR, high availability (HyperSwap), and the disaster recovery (Synchronous and Asynchronous mirroring) ones. It is used outside the CLI reference, where broader concepts can be explained.

## Creating a HyperSwap relation

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

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

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **vol** | Object name | Local volume to be replicated (the Master). | N | N/A |
| **create_slave** | Boolean | Determines whether to create a new Slave volume or to use an existing one. If an existing remote volume is used, its name must match the local volume name. | N | no |
| **remote_pool** | Object name | The storage pool on the remote system. Relevant only if creating a Slave volume. | N | N/A |

© Copyright IBM Corp. 2016, 2022

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **cg** | Object name | The local consistency group to be mirrored. | N | N/A |
| **slave_cg** | Object name | The name of the Slave consistency group on the remote storage system. | N | N/A |
| **target** | Object name | The remote target to contain the Slave volume. | Y | N/A |
| **init_type** | Enumeration | The initialization method of the Slave volume. | N | [none] |
| **part_of_multisite** | Boolean | Marks the HA as part of Multi-site. | N | no |

This command defines a new HyperSwap relation between Master and Slave peers. When you define a HyperSwap relation, the following rules are enforced:

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

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

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

The following initialization methods are available:

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

A storage system can have multiple HyperSwap relations between pairs of peers on various remote systems. However, when the peers have consistency group HyperSwap relations, all the volumes included in a specific consistency group HyperSwap relation can only be replicated between one pair of storage systems. Therefore, when a volume peer on a storage system (for example: A) has a HyperSwap relation with a volume on a remote storage system (for example: B), any other volume in the same consistency group on storage system A can only be defined in a HyperSwap relation with a volume on storage system B. The same is true for volumes from storage system B to A. In addition, the HyperSwap consistency group has one sync job for all pertinent HyperSwap volumes within the consistency group. To create a HyperSwap relation as part of a Multi-site relation, apply the **part_of_multisite** parameter.

**Example:**

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

**Output:**

```
Command completed successfully.
```

## Access control

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

## Warnings

- **VOLUME_SIZE_VERY_LARGE_ARE_YOU_SURE**

  The volume size is very large. It may not be possible to mirror this volume to older versions of the storage system. Are you sure?

## Return codes

- **BAD_REMOTE_VOLUME_NAME**

  The secondary volume name does not exist.
- **BAD_REMOTE_VOLUME_SIZE**

  The primary and secondary volumes contain a different number of blocks.
- **CONS_GROUP_BAD_NAME**

  The consistency group name does not exist.
- **CONS_GROUP_HAS_MIRROR**

  Mirroring is defined for this consistency group.
- **CONS_GROUP_MIRRORING_NOT_SUPPORTED_IN_TARGET**

  Consistency group mirroring is not supported by the target machine.
- **LOCAL_MAX_HA_REACHED**

  The maximum number of HyperSwap relationships is already reached on the local machine.
- **NOT_ENOUGH_SPACE_ON_REMOTE_MACHINE**

  Not enough free space to set the requested size of the secondary volume.
- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **VOLUME_IS_MASTER**

  This local volume is already defined as a primary volume.
- **VOLUME_IS_SLAVE**

  The volume is defined as a secondary volume.
- **VOLUME_HAS_OLVM**

  An IBM Hyper-Scale Mobility relationship is defined for this volume.
- **VOLUME_HAS_HA**

  This operation is forbidden on a volume with a HyperSwap relation.
- **TARGET_VOLUME_HAS_OLVM**

  This target volume is part of an IBM Hyper-Scale Mobility relationship.

- **TARGET_VOLUME_HAS_HA**

  This operation is forbidden, if the target volume is a peer in a HyperSwap relation.
- **REMOTE_VOLUME_EXISTS**

  The secondary volume with the indicated name already exists. The name cannot be reused.
- **REMOTE_MAX_VOLUMES_REACHED**

  The maximum number of volumes on the remote machine is already reached.
- **REMOTE_MAX_HA_REACHED**

  The maximum number of HyperSwap relationships is already reached on the remote machine.
- **REMOTE_CONS_GROUP_NOT_EMPTY**

  The remote consistency group contains volumes.
- **VOLUME_BAD_PREFIX**

  The volume name has a reserved prefix.
- **REMOTE_POOL_DOES_NOT_EXIST**

  The pool does not exist on the remote machine.
- **REMOTE_POOL_NOT_SPECIFIED**

  Prior to creating a secondary volume, a pool must be defined on the remote machine.
- **REMOTE_TARGET_NOT_CONNECTED**

  There is currently no connection from the target system.
- **VOLUME_IS_SNAPSHOT**

  The operation is not permitted on snapshots.
- **REMOTE_VOLUME_IS_SNAPSHOT**

  The secondary volume is a snapshot.
- **TARGET_BAD_NAME**

  The target name does not exist.
- **TARGET_BAD_TYPE**

  The target machine is not an XIV machine.
- **TARGET_NO_ACCESS**

  No access permissions to the secondary machine.
- **TARGET_NOT_CONNECTED**

  There is currently no connection to the target system.
- **REMOTE_VOLUME_LOCKED**

  The secondary volume is locked.
- **TIMEOUT**

  A remote operation was not completed in time.
- **VOLUME_HAS_MIRRORING_SNAPSHOTS**

  The volume has snapshots created by a previous mirroring process.
- **SLAVE_VOLUME_NOT_FORMATTED**

  The secondary volume is not formatted.
- **TARGET_DOES_NOT_ACCEPT_XIV_COMMANDS**

  The target system does not accept XIV management commands.
- **REMOTE_CONS_GROUP_IS_MIRRORED**

Mirroring is defined for this remote consistency group.

- **REMOTE_CONS_GROUP_BAD_NAME**

  The remote consistency group name does not exist.
- **REMOTE_VOLUME_IS_MASTER**

  A volume on the remote machine is already defined as primary.
- **REMOTE_VOLUME_IS_SLAVE**

  A volume on the remote machine is already defined as secondary.
- **REMOTE_MAX_MIRROR_CAPACITY_REACHED**

  The maximum capacity for mirrored volumes is already reached on the remote machine.
- **HA_RETRY_OPERATION**

  An operation is in progress on this HyperSwap relation.

  **Troubleshooting:** Try issuing the command again in a few seconds.
- **HA_INCOMPATIBLE_TARGET_VERSION**

  The automatic failover is not supported between the system versions of the specified peers.
- **NO_OFFLINE_INIT_TYPE_WITH_SLAVE_CREATION**

  A new volume will be created as secondary. Offline initialization is meaningless.
- **VOLUME_SIZE_ABOVE_LIMIT**

  The specified volume size is above the limit.
- **REMOTE_VOLUME_SIZE_ABOVE_LIMIT**

  The specified volume size is above the limit of the remote machine.
- **INVALID_SLICE_OFFSET**

  Slice offset is illegal.
- **VOLUME_IS_OLVM_PROXY**

  The volume is in an IBM Hyper-Scale Mobility Proxy phase.
- **REMOTE_VOLUME_IS_OLVM_PROXY**

  The remote volume is in an IBM Hyper-Scale Mobility Proxy phase.
- **ENCRYPTION_IN_PROGRESS**

  The system is in the process of changing the encryption activation state.
- **MIRROR_OF_SAME_TYPE_EXISTS_ON_VOLUME**

  A mirror of this type is already defined on this volume.
- **MIRROR_EXISTS_ON_TARGET**

  The volume already has a mirror on this target.
- **REMOTE_VOLUME_IS_MIRROR_MASTER**

  The volume is primary in a mirror relationship, and cannot be secondary!
- **MULTISITE_MAX_NUM_OF_MIRRORS_REACHED**

  Failed to create the relation, the maximum allowed number of relations is already exceeded.
- **REMOTE_VOLUME_TWO_SYNC_MIRRORS_NOT_ALLOWED**

  Two synchronous mirrors were detected on the remote volume. This is not allowed.
- **REMOTE_VOLUME_MIRROR_LOOP_DETECTED**

A mirror loop was detected on the remote volume. This means that there is a mirror on the remote system, whose target is this system. Therefore, you cannot create a mirror with this target on this system.

- **DOMAIN_MAX_MIRRORS_REACHED**

  The domain exceeds the maximum allowed number of mirrors.
- **REMOTE_DOMAIN_MAX_VOLUMES_REACHED**

  The maximum number of volumes in the remote machine domain is already reached.
- **REMOTE_DOMAIN_HAS_NO_ACCESS_TO_TARGET**

  The secondary machine domain has no access to the target.
- **DOMAIN_HAS_NO_ACCESS_TO_TARGET**

  The domain has no access to the target.
- **REMOTE_DOMAIN_MAX_MIRRORS_REACHED**

  The maximum number of mirrors is already reached in the remote machine domain.
- **DOMAIN_MAX_VOLUMES_REACHED**

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

  Data migration is already defined for the secondary volume.
- **REMOTE_VOLUME_MASTER_ASYNC_MIRROR_DETECTED**

  An asynchronous primary mirror was detected on the remote volume. The operation not allowed.
- **MAX_MULTISITES_REACHED**

  The number of Multi-site objects exceeded the limit.
- **REMOTE_VOLUME_HAS_MIRRORING_SNAPSHOTS**

  The remote volume has snapshots created by a previous mirroring process.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support
- **SYSTEM_OUT_OF_PHYSICAL_SPACE**

  The operation not allowed while the system is out of physical space.
- **REMOTE_DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier of the remote system is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support
- **REMOTE_SYSTEM_OUT_OF_PHYSICAL_SPACE**

  The operation not allowed while the remote system is out of physical space.
- **SLAVE_VOLUME_IS_MAPPED**

  The secondary volume is mapped.
- **TARGET_HAS_NO_QUORUM_WITNESS**

  The local target does not have a Quorum Witness defined.
- **REMOTE_TARGET_HAS_NO_QUORUM_WITNESS**

  The target on the remote system does not have a Quorum Witness defined.
- **HA_TARGET_QUORUM_WITNESS_IS_NOT_ACTIVATED**

  The Quorum Witness associated with the target is deactivated.

- **HA_REMOTE_TARGET_QUORUM_WITNESS_IS_NOT_ACTIVATED**

  The Quorum Witness associated with the remote target is deactivated.
- **HA_CONNECTIVITY_NOT_SUFFICIENT**

  The connectivity between the systems is not sufficient for the automatic failover.
- **HA_LOCAL_PEER_HAS_NO_QUORUM_WITNESS_CONNECTIVITY**

  The local peer connectivity to the Quorum Witness is not operational.
- **HA_REMOTE_PEER_HAS_NO_QUORUM_WITNESS_CONNECTIVITY**

  The remote peer connectivity to the Quorum Witness is not operational.
- **MAX_NUM_OF_PROXY_VOLUME_REACHED**

  Failed to create a mirror, because the maximum number of proxy volumes is exceeded.
- **REMOTE_MAX_NUM_OF_PROXY_VOLUME_REACHED**

  Failed to create a mirror, because the maximum number of remote proxy volumes is exceeded.
- **REMOTE_MAX_METADATA_OBJECTS_REACHED**

  The maximum number of metadata objects has been reached on a remote system.
- **HA_PEER_QUORUM_WITNESS_CONFIGURATION_NOT_VERIFIED**

  Connectivity between the peer and the Quorum Witness is not verified.
- **HA_REMOTE_PEER_QUORUM_WITNESS_CONFIGURATION_NOT_VERIFIED**

  Connectivity between the remote peer and the Quorum Witness is not verified.
- **TARGET_PEER_NOT_HEALTHY**

  The target peer is not identified as healthy.

  **Troubleshooting:** Check the Quorum Witness configuration.
- **HOST_TYPE_IS_NOT_CONFIGURED**

  Cannot associate a HyperSwap volume with a host of unconfigured type. IMPORTANT: Read the HyperSwap chapter in the 'Best Practices' document to understand the solution requirements.
- **CONS_GROUP_IS_SECONDARY**

  The consistency group is already defined as a secondary consistency group.
- **CONS_GROUP_IS_PRIMARY**

  The consistency group is already defined as a primary consistency group.
- **RELATION_OF_SAME_TYPE_EXISTS_ON_CONS_GROUP**

  A relation of this type is already defined on this consistency group.
- **MULTISITE_INCOMPATIBLE_TARGET_VERSION**

  Multi-site is not supported between the system versions of the specified peers.
- **REMOTE_VOLUME_BAD_POOL**

  The remote volume and remote consistency group belong to different storage pools.
- **REMOTE_VOLUME_BELONGS_TO_CONS_GROUP**

  The remote volume belongs to a consistency group.
- **CONS_GROUP_HA_ROLE_MISMATCH**

  All volumes in a HyperSwap consistency group must have the same mirroring role.
- **HA_HIGH_AVAILABILITY_ENABLED_IN_VOL**

  The consistency group's high availability is disabled but the volume's high availability is enabled.
- **HA_HIGH_AVAILABILITY_DISABLED_IN_VOL**

The consistency group's high availability is enabled, but the volume's high availability is disabled.

- **CONS_GROUP_RELATION_STATE_IS_NOT_INITIALIZING**

  This operation is only allowed when the matching consistency group relation state is 'Initializing'.
- **VOLUME_RELATION_IN_CG_CAN_NOT_BE_CREATED**

  A second relation must first be added to the consistency group before it can be added to a member volume.
- **CONS_GROUP_RELATION_PART_OF_MULTISITE_MISMATCH**

  All volumes in a mirrored or HyperSwap consistency group must have the same part_of_multisite setting.
- **CONS_GROUP_NOT_EMPTY**

  This operation is only allowed on an empty consistency group.
- **CONS_GROUP_HA_TARGET_MISMATCH**

  All volumes in a mirrored consistency group must have the same HyperSwap target.
- **CONS_GROUP_RELATION_INIT_TYPE_MISMATCH**

  All volumes in a mirrored or HyperSwap consistency group must have the same init type setting.
- **REMOTE_MULTISITE_CONS_GROUP_MEMBER_VOL_HAS_TOO_MANY_RELATIONS**

  A remote volume belonging to the consistency group has too many relations.
- **LOCAL_VOLUME_HAS_TOO_MANY_METADATA_OBJECTS**

  Local volume has too many metadata objects.

  **Troubleshooting:** Contact IBM Support.

# Viewing the status of HyperSwap volumes and consistency groups

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

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

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **scope** | Enumeration | List type: all HyperSwap mirrors, volumes, and consistency groups. | N | All (if no value is specified) |
| **vol** | Object name | Local volume name. | N | [none] |
| **cg** | Object name | Local consistency group name. | N | [none] |
| **target** | Object name | Remote target name. | N | [none] |
| **domain** | Object name | The domain name. | N | All Domains |

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

The following default parameters are shown:

- Name
- HA Object: Volume or CG
- Role: Master or Slave
- Remote System: The name of the target storage system

- Active: Yes or No
- Status: Initializing, Synchronized, or Unsynchronized
- Link Up: Yes or No
- Automatic Failover: The ability of a Slave volume to perform automatic failover: Active, Inactive, or N/A. Can be retrieved from either Master or Slave. N/A is returned for a volume that is part of a consistency group.

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

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

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

| Field ID | Field output | Description | Default position |
|---|---|---|---|
| **part_of_multisite** | part_of_multisite | N/A | N/A |
| **arch** | Remote Arch | N/A | N/A |
| **multisite_uid** | multisite_uid | N/A | N/A |

**Example:**

```
ha_list
```

**Output:**

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

### Access control

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

## Activating a HyperSwap relation

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

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

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **vol** | Object name | Master volume name.<br>If this parameter is not used, the **cg** parameter must be specified. | N | N/A |
| **cg** | Object name | Master consistency group name.<br>If this parameter is not used, the **vol** parameter must be specified. | N | N/A |
| **target** | Object name | Target HyperSwap relation name. | N | [none] |

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

The following is required for a successful command completion:

- The Master and Slave are configured on the same Quorum Witness
- The connectivity of Master and Slave to the Quorum Witness is healthy
- The specified object is a Master
- The connectivity between Master and Slave is sufficient for HyperSwap (that is, the **link_up** attribute of the **ha_list** command output is *yes*

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

A HyperSwap relation cannot be activated, if:

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

**Example:**

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

**Output:**

```
Command completed successfully.
```

## Access control

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

## Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **CONS_GROUP_BAD_NAME**

  The consistency group name does not exist.
- **LOCAL_PEER_IS_NOT_MASTER**

  The local peer is not primary.
- **HA_CONFIGURATION_ERROR**

  The HyperSwap relation's local configuration does not match its remote configuration.

- **REMOTE_MAX_VOLUMES_REACHED**

  The maximum number of volumes on the remote machine is already reached.
- **SYNC_ALREADY_ACTIVE**

  Synchronization is already active.
- **VOLUME_BELONGS_TO_HA_CONS_GROUP**

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

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

  The automatic failover cannot be activated.
- **HA_CONS_GROUP_MEMBERSHIP_MISMATCH**

  The HyperSwap consistency group contains different primary and secondary volumes. This may have happened because the cg_add_vol or cg_remove_vol command was previously issued, but the primary did not receive an acknowledgment from the secondary until the command timed out, or due to any other unexpected failure.
- **HA_SIZE_MISMATCH**

  The sizes of the primary and secondary volumes in this HyperSwap relation are different.
- **HA_RETRY_OPERATION**

  An operation is in progress on this HyperSwap relation.

  **Troubleshooting:** Try issuing the command again in a few seconds.
- **TARGET_BAD_NAME**

  The target name does not exist.
- **VOLUME_TARGET_MISMATCH**

  The volume and target do not match.
- **CONS_GROUP_BAD_TARGET**

  The target name does not match the consistency group.
- **VOLUME_TOO_MANY_ACTIVE_MIRRORS**

  This command cannot be issued if more than one mirror is active on the volume.
- **REMOTE_DOMAIN_MAX_VOLUMES_REACHED**

  The maximum number of volumes in the remote machine domain is already reached.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support
- **REMOTE_DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier of the remote system is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support
- **REMOTE_SYSTEM_OUT_OF_PHYSICAL_SPACE**

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

  The target peer is not identified as healthy.

  **Troubleshooting:** Check the Quorum Witness configuration.

- **TARGET_HAS_NO_QUORUM_WITNESS**

  The local target does not have a Quorum Witness defined.
- **REMOTE_TARGET_HAS_NO_QUORUM_WITNESS**

  The target on the remote system does not have a Quorum Witness defined.
- **VOLUME_IS_NOT_HA**

  This is not a HyperSwap volume.
- **CONS_GROUP_IS_NOT_HA**

  The local consistency group does not have HyperSwap definitions.
- **HA_CONNECTIVITY_NOT_SUFFICIENT**

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

  Connectivity between the peer and the Quorum Witness is not verified.
- **HA_REMOTE_PEER_QUORUM_WITNESS_CONFIGURATION_NOT_VERIFIED**

  Connectivity between the remote peer and the Quorum Witness is not verified.
- **HA_TARGET_QUORUM_WITNESS_IS_NOT_ACTIVATED**

  The Quorum Witness associated with the target is deactivated.
- **HA_REMOTE_TARGET_QUORUM_WITNESS_IS_NOT_ACTIVATED**

  The Quorum Witness associated with the remote target is deactivated.
- **CONS_GROUP_MEMBER_VOL_IS_MISSING_A_RELATION**

  A volume in the consistency group is missing the volume-level Multi-site/mirror/HyperSwap relation.
- **REMOTE_MULTISITE_IS_NOT_SMASTER**

  The remote peer is not the SMaster.
- **MULTISITE_SMASTER_INVALID_CONFIGURATION**

  The Multi-site SMaster-Master relation configuration is invalid.
- **MULTISITE_BAD_GLOBAL_ID**

  The Multi-site global ID does not exist.

## Deactivating a HyperSwap relation

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

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

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **vol** | Object name | Master volume name or a list of master volumes.<br>If this paramater is not used, the **cg** parameter must be specified. | N | N/A |

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **cg** | Object name | Master consistency group name or a list of master consistency groups.<br>If this parameter is not used, the **vol** parameter must be specified. | N | N/A |
| **target** | Object name | Target HyperSwap relation name. | N | [none] |

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

**Example:**

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

**Output:**

```
Command completed successfully.
```

## Access control

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

## Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **VOLUME_IS_NOT_HA**

  This is not a HyperSwap volume.
- **CONS_GROUP_BAD_NAME**

  The consistency group name does not exist.
- **CONS_GROUP_IS_NOT_HA**

  The local consistency group does not have HyperSwap definitions.
- **LOCAL_PEER_IS_NOT_MASTER**

  The local peer is not primary.
- **SYNC_ALREADY_INACTIVE**

  Synchronization is already inactive.
- **VOLUME_BELONGS_TO_HA_CONS_GROUP**

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

- **HA_RETRY_OPERATION**

  An operation is in progress on this HyperSwap relation.

  **Troubleshooting:** Try issuing the command again in a few seconds.
- **TARGET_BAD_NAME**

  The target name does not exist.
- **VOLUME_TARGET_MISMATCH**

  The volume and target do not match.
- **CONS_GROUP_BAD_TARGET**

  The target name does not match the consistency group.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support

# Deleting a HyperSwap relation

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

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

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **vol** | Object name | Local Master volume name.<br>If this parameter is not used, the **cg** parameter must be specified. | N | N/A |
| **cg** | Object name | Local Master consistency group name.<br>If this parameter is not used, the **vol** parameter must be specified. | N | N/A |
| **target** | Object name | Target HyperSwap relation name. | N | [none] |
| **force_on_slave** | Boolean | Forces the deletion of the HyperSwap relation on the Slave peer. Available only if the Slave is in the Initializing state. | N | no |

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

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

- The HyperSwap relation is in the Standby state, in other words, it has just been created or deactivated.
- The communication is established. If there is no communication, the HyperSwap relation is only deleted on the Master. When the communication resumes, a configuration error will be issued on the Slave.

- The Slave volume is not mapped.

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

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

**Deletion when the HyperSwap relation is inactive or when the connectivity has failed**

**On an active Master**

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

**On an active Slave**

**Important:**

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

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

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

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

**On a disconnected Master or Slave**

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

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

**Example:**

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

**Output:**

```
Command completed successfully.
```

## Access control

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

## Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_DELETE_CG_HA**

  Are you sure you want to delete the HyperSwap relations of the consistency group and of all the volumes in it?

## Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **VOLUME_IS_NOT_HA**

  This is not a HyperSwap volume.
- **CONS_GROUP_IS_NOT_HA**

  The local consistency group does not have HyperSwap definitions.
- **CONS_GROUP_BAD_NAME**

  The consistency group name does not exist.
- **LOCAL_PEER_IS_NOT_MASTER**

  The local peer is not primary.
- **HA_IS_ACTIVE**

  HyperSwap is currently active.
- **FORCE_DELETE_NOT_ALLOWED_ON_MASTER**

  Deletion needs to be forced on secondary mirrors only.
- **VOLUME_BELONGS_TO_HA_CONS_GROUP**

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

  An operation is in progress on this HyperSwap relation.

  **Troubleshooting:** Try issuing the command again in a few seconds.
- **HA_IS_NOT_INITIALIZING**

  Deleting a HyperSwap relation is permitted only during the initialization phase.
- **TARGET_BAD_NAME**

  The target name does not exist.
- **VOLUME_TARGET_MISMATCH**

  The volume and target do not match.
- **CONS_GROUP_BAD_TARGET**

  The target name does not match the consistency group.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support
- **SLAVE_VOLUME_IS_MAPPED**

  The secondary volume is mapped.
- **MIRROR_CONS_GROUP_MEMBERSHIP_MISMATCH**

  The mirrored consistency group contains different volumes on the primary and secondary machines. This problem occurs whenever the cg_add_vol or cg_remove_vol commands were previously issued, and the primary machine did not receive an acknowledgment from the secondary machine until the command timed out, or due to any other unexpected failure.
- **VOLUME_IS_AN_UNAVAILABLE_HYPERSWAP_PEER**

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

This HyperSwap relation is associated with a defined Multi-site, the operation is not allowed.

- **HA_PART_OF_MULTISITE**

  This HyperSwap relation is part of a Multi-site relation.

# Switching roles between Master and Slave volumes

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

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

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **vol** | Object name | Local volume name.<br>If this parameter is not used, the **cg** parameter must be specified. | N | N/A |
| **cg** | Object name | Local consistency group name.<br>If this parameter is not used, the **vol** parameter must be specified. | N | N/A |
| **target** | Object name | N/A | N | [none] |

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

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

Following the successful completion of the command:

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

**Example:**

```
ha_switch_roles vol=DBVolume target=SecondarySite
```

**Output:**

```
Command completed successfully
```

## Access control

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

## Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_SWITCH_ROLES**

  Are you sure you want to switch the roles in this relation?

- **ARE_YOU_SURE_YOU_WANT_TO_SWITCH_THE_PEER_ROLES**

  The system that owns the primary volume is currently not connected to the Quorum Witness. If you switch the peer roles, the automatic failover may become unavailable. Are you sure you want to continue?

## Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.

- **CONS_GROUP_BAD_NAME**

  The consistency group name does not exist.

- **LOCAL_PEER_IS_NOT_MASTER**

  The local peer is not primary.

- **HA_IS_NOT_SYNCHRONIZED**

  The HyperSwap relation is not synchronized.

- **VOLUME_HAS_DATA_MIGRATION**

  Data Migration is defined for this volume.

- **REMOTE_TARGET_NOT_CONNECTED**

  There is currently no connection from the target system.

- **VOLUME_BELONGS_TO_HA_CONS_GROUP**

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

- **HA_CONNECTIVITY_NOT_SUFFICIENT**

  The connectivity between the systems is not sufficient for the automatic failover.

- **HA_RELATION_MASTER_COULD_NOT_UPDATE_QW_AFTER_RETURN_TO_GOOD_STATE**

  The primary volume had control over the relationship during a past failure, and was unable to update the Quorum Witness after the recovery.

- **HA_HAS_SYNC_JOB**

  This operation is not permitted on a HyperSwap relation with active sync jobs.

- **HA_RETRY_OPERATION**

  An operation is in progress on this HyperSwap relation.

  **Troubleshooting:** Try issuing the command again in a few seconds.

- **REMOTE_HA_IS_NOT_ACTIVE**

  The remote peer in this HyperSwap relation is not active.

- **TARGET_BAD_NAME**

  The target name does not exist.

- **VOLUME_TARGET_MISMATCH**

  The volume and target do not match.

- **CONS_GROUP_BAD_TARGET**

  The target name does not match the consistency group.

- **COMMAND_NOT_SUPPORTED_FOR_OLVM_VOLUMES**

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

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support
- **SYSTEM_OUT_OF_PHYSICAL_SPACE**

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

  This is not a HyperSwap volume.
- **CONS_GROUP_IS_NOT_HA**

  The local consistency group does not have HyperSwap definitions.
- **REMOTE_MAY_NOT_HAVE_COMPLETED_THE_OPERATION**

  The operation may be not yet completed on the remote target.
- **VOLUME_HAS_MULTIPLE_RELATIONS**

  The volume has multiple relations. The operation is not allowed or a target must be specified.
- **CONS_GROUP_HAS_MULTIPLE_RELATIONS**

  The consistency group has multiple relations. The operation is not allowed or a target must be specified.

# Changing a peer role in a HyperSwap volume

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

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

## Parameters

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

This command changes the role of the local peer from Master to Slave or from Slave to Master when the HyperSwap relation is not activated. The command should be issued on both peers before the relation is activated again, so that upon reconnection there still will be one Master and one Slave.

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

For a successful role change from Master to Slave, the volume can be in any phase, except *Initializing*. The Master ceases serving host requests, and is set to accept replication from the other peer as a Slave.

To successfully change a Slave to a Master, the HyperSwap relation must be deactivated.

**Before changing a Slave to a Master:**

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

**Warning:**

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

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

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

**Example:**

```
ha_change_role vol=regular_volume
```

**Output:**

```
Command completed successfully.
```

## Access control

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

## Warnings

- **SOME_DATA_WILL_BE_LOST_ARE_YOU_SURE**

  Are you sure you want the mirror/HyperSwap local peer to become secondary and lose the data that was not replicated?

- **ARE_YOU_SURE_YOU_WANT_TO_CHANGE_A_PEER_WITH_LCS_TO_MASTER**

  Are you sure you want the mirror/HyperSwap local peer to become primary? The local peer has a last-consistent snapshot.

## Return codes

- **VOLUME_IS_NOT_HA**

  This is not a HyperSwap volume.
- **CONS_GROUP_IS_NOT_HA**

  The local consistency group does not have HyperSwap definitions.
- **HA_IS_ACTIVE**

  HyperSwap is currently active.
- **HA_IS_INITIAL**

  The operation is not permitted during the HyperSwap relation initialization phase.
- **VOLUME_BELONGS_TO_HA_CONS_GROUP**

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

  The volume name does not exist.
- **CONS_GROUP_BAD_NAME**

  The consistency group name does not exist.
- **VOLUME_HAS_DATA_MIGRATION**

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

  An operation is in progress on this HyperSwap relation.

  **Troubleshooting:** Try issuing the command again in a few seconds.
- **HA_HAS_NO_SYNCHED_SNAPSHOT**

  This HyperSwap volume does not have a synchronized snapshot.
- **MASTER_CANNOT_BE_DEMOTED**

  The primary volume cannot be demoted to secondary. Peer status mismatch.
- **TARGET_BAD_NAME**

  The target name does not exist.
- **VOLUME_TARGET_MISMATCH**

  The volume and target do not match.
- **CONS_GROUP_BAD_TARGET**

  The target name does not match the consistency group.
- **COMMAND_NOT_SUPPORTED_FOR_OLVM_VOLUMES**

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

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support
- **SYSTEM_OUT_OF_PHYSICAL_SPACE**

The operation not allowed while the system is out of physical space.

- **VOLUME_HAS_MULTIPLE_RELATIONS**

  The volume has multiple relations. The operation is not allowed or a target must be specified.

- **CONS_GROUP_HAS_MULTIPLE_RELATIONS**

  The consistency group has multiple relations. The operation is not allowed or a target must be specified.

- **HA_PART_OF_MULTISITE**

  This HyperSwap relation is part of a Multi-site relation.

## Restoring the availability of a Master volume

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

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

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **vol** | Object name | Local volume name.<br>Must be specified if the command is applied to a volume. | N |
| **cg** | Object name | CG name<br>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 |

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

### Return codes

- **VOLUME_IS_NOT_HA**

  This is not a HyperSwap volume.
- **CONS_GROUP_IS_NOT_HA**

  The local consistency group does not have HyperSwap definitions.
- **CONS_GROUP_BAD_NAME**

  The consistency group name does not exist.
- **HA_IS_ACTIVE**

  HyperSwap is currently active.
- **VOLUME_BELONGS_TO_HA_CONS_GROUP**

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

  The volume name does not exist.
- **LOCAL_PEER_IS_NOT_MASTER**

  The local peer is not primary.

## Creating a HyperSwap volume snapshot (ad hoc sync job)

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

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

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **vol** | Object name | The name of the volume to create a snapshot for. | N | N/A |
| **cg** | Object name | Local master consistency group name. | N | N/A |
| **target** | Object name | Target HyperSwap relation name. Mandatory if 2 mirrors are defined on the volume. | N | [none] |
| **name** | Object name | The name of the new snapshot. | Y | N/A |
| **slave_name** | Object name | The name of the new snapshot on the slave. | N | [none] |

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **delete_priority** | Integer | The deletion priority of the volume's snapshot. | N | 1 |

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

Prerequisite:

- The HyperSwap relation is operational and synchronized.

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

**Example:**

```
ha_create_snapshot vol=simpleVolume target=myTarget name=mySnapshot slave_name=slave1
delete_priority=3
```

**Output:**

```
Command completed successfully
```

## Access control

| User Category | Permission | Condition |
|---|---|---|
| Storage administrator | Allowed | N/A |
| Storage integration administrator | Allowed | N/A |
| Application administrator | Conditionally Allowed | The volume is mapped to a host or a cluster associated with the user. If a snapshot overwrite is used, the target snapshot must be one created by a server administrator. |
| Security administrator | Disallowed | N/A |
| Read-only users | Disallowed | N/A |
| Technicians | Disallowed | N/A |

## Return codes

- **CONS_GROUP_MISMATCH**

  The snapshot group does not match the consistency group volumes.
- **CONS_GROUP_EMPTY**

  The operation is not allowed on an empty consistency group.
- **CONS_GROUP_BAD_NAME**

  The consistency group name does not exist.
- **VOLUME_IS_NOT_HA**

  This is not a HyperSwap volume.
- **CONS_GROUP_IS_NOT_HA**

  The local consistency group does not have HyperSwap definitions.
- **LOCAL_PEER_IS_NOT_MASTER**

  The local peer is not primary.
- **HA_IS_NOT_SYNCHRONIZED**

The HyperSwap relation is not synchronized.

- **HA_RETRY_OPERATION**

  An operation is in progress on this HyperSwap relation.

  **Troubleshooting:** Try issuing the command again in a few seconds.
- **HA_IS_NON_OPERATIONAL**

  This HyperSwap volume is not operational.
- **MAX_VOLUMES_REACHED**

  The maximum allowed number of volumes is already reached.
- **DOMAIN_MAX_VOLUMES_REACHED**

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

  The requested operation is not allowed on a loopback target.
- **OVERWRITE_SNAPSHOT_BAD_NAME**

  The snapshot name does not exist.
- **OVERWRITE_SNAPSHOT_GROUP_DOES_NOT_BELONG_TO_GIVEN_GROUP**

  The snapshot group belongs to another consistency group.
- **POOL_SNAPSHOT_LIMIT_REACHED**

  There is not enough space to create a snapshot.
- **REMOTE_POOL_SNAPSHOT_LIMIT_REACHED**

  There is not enough space on the remote target for creating a snapshot.
- **REMOTE_MAX_VOLUMES_REACHED**

  The maximum number of volumes on the remote machine is already reached.
- **REMOTE_MAX_SNAPSHOTS_FOR_VOLUME_REACHED**

  The maximum allowed number of snapshots per volume is already reached on a remote machine whose version is not 10.2.4.
- **REMOTE_VOLUME_IS_MASTER**

  A volume on the remote machine is already defined as primary.
- **REMOTE_VOLUME_IS_SNAPSHOT**

  The secondary volume is a snapshot.
- **REMOTE_SNAPSHOT_NAME_EXISTS**

  The remote snapshot name already exists.
- **REMOTE_SNAPSHOT_ILLEGAL_PRIORITY**

  Illegal snapshot priority (remote); must be an integer between 1 and 4.
- **REMOTE_SNAPSHOT_GROUP_NAME_EXISTS**

  The remote snapshot group name already exists.
- **REMOTE_SNAPSHOT_GROUP_ILLEGAL_PRIORITY**

  Illegal snapshot group priority (remote); must be an integer between 1 and 4.
- **REMOTE_SNAPSHOT_GROUP_BAD_PREFIX**

  The remote snapshot group name has a reserved prefix.
- **REMOTE_SNAPSHOT_BAD_PREFIX**

  The remote snapshot name has a reserved prefix.

- **REMOTE_CONS_GROUP_EMPTY**

  The operation is not allowed on an empty consistency group (remote).
- **REMOTE_CONS_GROUP_MISMATCH**

  The remote snapshot group does not match the consistency group volumes.
- **SNAPSHOT_ILLEGAL_PRIORITY**

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

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

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

  The snapshot group name already exists.
- **SNAPSHOT_GROUP_ILLEGAL_PRIORITY**

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

  The snapshot group name does not exist.
- **SNAPSHOT_GROUP_BAD_PREFIX**

  The snapshot group name has a reserved prefix.
- **SNAPSHOT_IS_PART_OF_SNAPSHOT_GROUP**

  The snapshot is part of a snapshot group.
- **VOLUME_BAD_PREFIX**

  The volume name has a reserved prefix.
- **VOLUME_BELONGS_TO_HA_CONS_GROUP**

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

  The volume name already exists.
- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **VOLUME_IS_NOT_CONSISTENT_SLAVE**

  The operation not allowed on an inconsistent secondary volume.
- **VOLUME_IS_SNAPSHOT**

  The operation is not permitted on snapshots.
- **CONS_GROUP_BAD_TARGET**

  The target name does not match the consistency group.
- **TARGET_BAD_NAME**

  The target name does not exist.
- **VOLUME_TARGET_MISMATCH**

  The volume and target do not match.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

**Troubleshooting:** Contact IBM Support

- **MAX_SNAPSHOTS_PER_VOLUME_REACHED**

  The maximum allowed number of snapshots is already reached.
- **REMOTE_MAX_SNAPSHOTS_PER_VOLUME_REACHED**

  The maximum allowed number of snapshots is already reached on the remote ssytem.
- **REMOTE_DOMAIN_MAX_VOLUMES_REACHED**

  The maximum number of volumes in the remote machine domain is already reached.
- **REMOTE_CONS_GROUP_BAD_NAME**

  The remote consistency group name does not exist.
- **SNAPSHOT_CAN_NOT_BE_CREATED_REMOTE_CONS_GROUP_IO_IS_NOT_PAUSED**

  The snapshot group will not be created since the remote consistency group is not in a stopped state.
- **SNAPSHOT_CAN_NOT_BE_CREATED_REMOTE_CONS_GROUP_DEFINITION_CHANGED**

  The snapshot group will not be created since the volumes in the remote consistency group have changed since the io_pause command was issued.
- **REMOTE_OVERWRITE_SNAPSHOT_GROUP_DOES_NOT_BELONG_TO_GIVEN_GROUP**

  The remote snapshot group belongs to another consistency group.

# Changing the designation of HyperSwap relation peers

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

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

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **vol** | Object name | Master volume name. | N | N/A |
| **cg** | Object name | Master consistency group name. | N | N/A |
| **target** | Object name | The name of the target. | N | [none] |
| **new_designation** | Enumeration | The new designation of the peer<br>If not specified, the command swaps the designation of the Primary and Secondary peers. | N | none |

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

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

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

**Example:**

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

**Output:**

```
Command completed successfully.
```

## Access control

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

## Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **VOLUME_IS_NOT_HA**

  This is not a HyperSwap volume.
- **CONS_GROUP_BAD_NAME**

  The consistency group name does not exist.
- **CONS_GROUP_IS_NOT_HA**

  The local consistency group does not have HyperSwap definitions.
- **LOCAL_PEER_IS_NOT_MASTER**

  The local peer is not primary.
- **MIRROR_DESIGNATION_NOT_SUPPORTED_BY_TARGET**

  The mirror's target does not support mirror role designation.
- **HA_IS_NON_OPERATIONAL**

  This HyperSwap volume is not operational.
- **VOLUME_BELONGS_TO_HA_CONS_GROUP**

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

  The target name does not exist.
- **VOLUME_TARGET_MISMATCH**

  The volume and target do not match.
- **CONS_GROUP_BAD_TARGET**

  The target name does not match the consistency group.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support

# Enabling high availability of a HyperSwap relation

Use the **ha_high_availability_enable** command to enable high availability (automatic failover) of a HyperSwap relation.

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

## Parameters

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

The high availability is enabled in order to allow automatic failover, after the high availability had been disabled in the past. The actual state of the automatic failover, as reported in **ha_list**, depends on several factors, including the user enablement and data synchronization state. It is possible that even though the high availability is enabled by this command, the automatic failover remains inactive due to other factors.

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

**Example:**

```
ha_high_availability_enable vol=regular_volume
```

**Output:**

```
Command completed successfully.
```

## Access control

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

## Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **CONS_GROUP_BAD_NAME**

  The consistency group name does not exist.
- **VOLUME_IS_NOT_HA**

  This is not a HyperSwap volume.
- **CONS_GROUP_IS_NOT_HA**

The local consistency group does not have HyperSwap definitions.

- **TARGET_NOT_CONNECTED**

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

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

  The local peer is not primary.
- **VOLUME_BELONGS_TO_HA_CONS_GROUP**

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

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

  Automatic failover is already enabled (valid only for ha_high_availability_enable).
- **HA_PEER_QUORUM_WITNESS_CONFIGURATION_NOT_VERIFIED**

  Connectivity between the peer and the Quorum Witness is not verified.
- **HA_REMOTE_PEER_QUORUM_WITNESS_CONFIGURATION_NOT_VERIFIED**

  Connectivity between the remote peer and the Quorum Witness is not verified.

## Disabling high availability of a HyperSwap relation

Use the **ha_high_availability_disable** command to disable the high availability (automatic failover) of a HyperSwap relation.

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

### Parameters

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

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

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

**Example:**

```
ha_high_availability_disable vol=regular_volume
```

**Output:**

```
Command completed successfully.
```

### Access control

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

### Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **CONS_GROUP_BAD_NAME**

  The consistency group name does not exist.
- **VOLUME_IS_NOT_HA**

  This is not a HyperSwap volume.
- **CONS_GROUP_IS_NOT_HA**

  The local consistency group does not have HyperSwap definitions.
- **TARGET_NOT_CONNECTED**

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

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

  The local peer is not primary.
- **VOLUME_BELONGS_TO_HA_CONS_GROUP**

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

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

  Automatic failover is already disabled (valid only for ha_high_availability_disable).

## Converting a HyperSwap relation into a sync mirror

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

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

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **vol** | Object name | Master volume. | N |

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **cg** | Object name | Master consistency group name or a list of master consistency groups. | N |

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

Prerequisites:

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

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

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

**Example:**

```
ha_convert_into_mirror vol=regular_volume
```

**Output:**

```
Command completed successfully.
```

## Access control

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

## Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **CONS_GROUP_BAD_NAME**

  The consistency group name does not exist.
- **VOLUME_IS_NOT_HA**

  This is not a HyperSwap volume.
- **CONS_GROUP_IS_NOT_HA**

  The local consistency group does not have HyperSwap definitions.
- **TARGET_NOT_CONNECTED**

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

There is currently no connection from the target system.

- **LOCAL_PEER_IS_NOT_MASTER**

  The local peer is not primary.

- **SLAVE_VOLUME_IS_MAPPED**

  The secondary volume is mapped.

- **VOLUME_BELONGS_TO_HA_CONS_GROUP**

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

- **HA_IS_NOT_OPERATIONAL**

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

- **VOLUME_HAS_MULTIPLE_RELATIONS**

  The volume has multiple relations. The operation is not allowed or a target must be specified.

- **CONS_GROUP_HAS_MULTIPLE_RELATIONS**

  The consistency group has multiple relations. The operation is not allowed or a target must be specified.

## Converting a sync mirror into a HyperSwap relation

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

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

### Parameters

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

Prerequisites

- The remote volume is not mapped
- The mirror is a sync mirror
- The local mirror is the Master
- The target connectivity is active and sufficient for HyperSwap (verified with the target_list command)
- The peers are connected to a properly configured and active Quorum Witness
- Both volumes have the same name

The command is issued on the Master peer and affects both peers, that are converted to HyperSwap at the same time. The HyperSwap functionality requires certain configuration elements, such as Quorum Witness, and will fail if they do not exist or are not in the required state.

Once the command has successfully completed, the relation will no longer be retrieved by the mirror_list command, but rather by the ha_list command.

The conversion process usually continues after the command completion event has been issued. While it continues, the indication of the high availability activation state Automatic Failover is Inactive, because metadata is still not synchronized. Once the conversion has completed, the HyperSwap volume can be mapped from the remote system and the new mapping can be used to perform I/O to that volume.

**Example:**

```
mirror_convert_into_ha vol=regular_volume
```

**Output:**

```
Command completed successfully.
```

## Access control

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

## Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **BAD_REMOTE_VOLUME_NAME**

  The secondary volume name does not exist.
- **CONS_GROUP_BAD_NAME**

  The consistency group name does not exist.
- **VOLUME_NO_MIRROR**

  The local volume does not have remote mirroring definitions.
- **CONS_GROUP_NO_MIRROR**

  The local consistency group does not have remote mirroring definitions.
- **MIRROR_IS_STANDBY**

  The mirror is marked as Standby.
- **TARGET_NOT_CONNECTED**

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

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

  The local peer is not primary.
- **SLAVE_VOLUME_IS_MAPPED**

  The secondary volume is mapped.
- **VOLUME_BELONGS_TO_MIRRORED_CONS_GROUP**

  The volume mirror is part of a consistency group mirror.
- **TARGET_HAS_NO_QUORUM_WITNESS**

  The local target does not have a Quorum Witness defined.
- **REMOTE_TARGET_HAS_NO_QUORUM_WITNESS**

  The target on the remote system does not have a Quorum Witness defined.

- **HA_LOCAL_PEER_HAS_NO_QUORUM_WITNESS_CONNECTIVITY**

  The local peer connectivity to the Quorum Witness is not operational.
- **HA_INCOMPATIBLE_TARGET_VERSION**

  The automatic failover is not supported between the system versions of the specified peers.
- **MIRROR_TYPE_IS_NOT_SYNC**

  Mirror type is not Sync Best Effort.
- **MIRROR_IS_NON_OPERATIONAL**

  The mirror is non-operational.
- **MIRROR_RETRY_OPERATION**

  There is an operation in progress on this mirror.

  **Troubleshooting:** Retry the command in a few seconds.
- **REMOTE_MAX_HA_REACHED**

  The maximum number of HyperSwap relationships is already reached on the remote machine.
- **LOCAL_MAX_HA_REACHED**

  The maximum number of HyperSwap relationships is already reached on the local machine.
- **VOLUME_HAS_OLVM**

  An IBM Hyper-Scale Mobility relationship is defined for this volume.
- **VOLUME_HAS_HA**

  This operation is forbidden on a volume with a HyperSwap relation.
- **MAX_NUM_OF_PROXY_VOLUME_REACHED**

  Failed to create a mirror, because the maximum number of proxy volumes is exceeded.
- **TARGET_VOLUME_HAS_OLVM**

  This target volume is part of an IBM Hyper-Scale Mobility relationship.
- **TARGET_VOLUME_HAS_HA**

  This operation is forbidden, if the target volume is a peer in a HyperSwap relation.
- **REMOTE_MAX_METADATA_OBJECTS_REACHED**

  The maximum number of metadata objects has been reached on a remote system.
- **LOCAL_AND_REMOTE_VOLUME_NAMES_ARE_DIFFERENT**

  Local and remote volume names are different.
- **HA_TARGET_QUORUM_WITNESS_IS_NOT_ACTIVATED**

  The Quorum Witness associated with the target is deactivated.
- **TARGET_PEER_NOT_HEALTHY**

  The target peer is not identified as healthy.

  **Troubleshooting:** Check the Quorum Witness configuration.
- **HA_CONNECTIVITY_NOT_SUFFICIENT**

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

  Connectivity between the peer and the Quorum Witness is not verified.
- **HA_REMOTE_PEER_QUORUM_WITNESS_CONFIGURATION_NOT_VERIFIED**

  Connectivity between the remote peer and the Quorum Witness is not verified.
- **HOST_TYPE_IS_NOT_CONFIGURED**

Cannot associate a HyperSwap volume with a host of unconfigured type. IMPORTANT: Read the HyperSwap chapter in the 'Best Practices' document to understand the solution requirements.

- **VOLUME_HAS_MULTIPLE_RELATIONS**

  The volume has multiple relations. The operation is not allowed or a target must be specified.

- **CONS_GROUP_HAS_MULTIPLE_RELATIONS**

  The consistency group has multiple relations. The operation is not allowed or a target must be specified.

- **LOCAL_VOLUME_HAS_TOO_MANY_METADATA_OBJECTS**

  Local volume has too many metadata objects.

  **Troubleshooting:** Contact IBM Support.

# Creating a new Quorum Witness

Use the **quorum_witness_define** command to create a new Quorum Witness definition in the system and connect the system to the Quorum Witness.

```
quorum_witness_define name=qw_name certificate=qw_certificate address=qw_address [ port=qw_port ]
 [ activate=<yes|no> ]
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **name** | Object name | The name of the Quorum Witness to be created. | Y | N/A |
| **certificate** | N/A | The public certificate or certificate chain of the Quorum Witness (see below for details). | Y | N/A |
| **address** | N/A | The Quorum Witness address: IPv4, IPv6 (full format only) or DNS name. | Y | N/A |
| **port** | Positive integer | The port used for Quorum Witness communications. | N | 8460 |
| **activate** | Boolean | Defines whether to activate the Quorum Witness upon creation. | N | yes |

This command defines a Quorum Witness to be used for HyperSwap relations. Up to two Quorum Witness instances may be defined. Setting the **activate** parameter to yes will also activate the Quorum Witness and start the communication between the system and the Quorum Witness.

**The certificate parameter:**

The value of the **certificate** parameter is the content of a PEM file with asterisks instead of newlines. Chained certificates are supported. The total maximal length of a PEM file holding chained certificates (leaf first, root last) is 15360 characters (including the asterisk characters). In Windows, you can drag-and-drop a PEM file from the Windows Explorer to the appropriate location in the CLI session window; the content will be added automatically.

**Example:**

```
quorum_witness_define name=q1 certificate="Certificate" address=192.0.2.1
```

**Output:**

```
Command completed successfully.
```

## Access control

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

## Return codes

- **QUORUM_WITNESS_NAME_ALREADY_EXISTS**

  The Quorum Witness name already exists.
- **QUORUM_WITNESS_ADDRESS_ALREADY_EXISTS**

  The Quorum Witness IPv4/6 address or DNS name with same port number already exists.

  **Troubleshooting:** Use the already configured Quorum Witness or change the address or port number.
- **QUORUM_WITNESS_MISSING_SERVICE_CERTIFICATE**

  No valid certificate is defined for the Quorum Witness service.

  **Troubleshooting:** Use PKI commands to define the certificate for the Quorum Witness service.
- **MAX_QUORUM_WITNESSES_REACHED**

  Too many Quorum Witness instances are defined.

  **Troubleshooting:** Delete one or more Quorum Witness instances and then try again.
- **MAX_ACTIVE_QUORUM_WITNESSES_REACHED**

  The maximum number of active Quorum Witness instances is already reached.

  **Troubleshooting:** Deactivate an active Quorum Witness and then try again.
- **SSL_CERTIFICATE_HAS_EXPIRED**

  The SSL certificate has expired.
- **SSL_CERTIFICATE_ISSUER_NOT_FOUND**

  The SSL certificate issuer was not found in the certificate chain.
- **SSL_CERTIFICATE_INVALID_FORMAT**

  The SSL certificate format is invalid or corrupted.
- **SSL_CERTIFICATE_CHAIN_EMPTY**

  No certificates were found in the input.
- **SSL_CERTIFICATE_VERIFICATION_INTERNAL_ERROR**

  The SSL certificate verification has failed because of an internal system error.
- **SSL_CERTIFICATE_VERIFICATION_FAILED**

  The SSL certificate chain verification failed.
- **SSL_CERTIFICATE_NOT_YET_VALID**

  The SSL certificate is not yet valid.

# Listing Quorum Witnesses

Use the **quorum_witness_list** command to list all Quorum Witnesses defined in the system, or only the specified one.

```
quorum_witness_list [ name=qw_name ] [ domain=DomainName ]
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **name** | Object name | The internal name of the Quorum witness. | N | All Quorum Witnesses |
| **domain** | Object name | The domain name. | N | All domains |

The command output includes the following fields:

- **name**: The Quorum Witness internal name
- **quorum_id**: A globally unique Quorum Witness ID
- **address**: The v6 or DNS address for communicating with the Quorum Witness
- **port**: The port for communicating with the Quorum Witness
- **state**: The state of the Quorum Witness in the system. The following values are available: `Activating`, `Activated`, `Deactivating`, and `Deactivated`.
- **connection**: The state of the connection with the Quorum Witness, accumulated across all Quorum Node connection statuses.
- **external_name**: The Quorum Witness external name
- **db_health**: The state of the Quorum Witness DB health. The following values are available:

| Value | Meaning |
|---|---|
| OK | N/A |
| `Recovery needed` | A problem was identified in the Quorum Witness DB, that prevents it from operating properly. |
| `Restore pending` | The Quorum Witness DB was initialized. Once the Quorum Witness is activated, the system will automatically re-register to the Quorum Witness, and restore the information pertaining to its HyperSwap relations. |

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

- **heartbeats_ok**: Indicates whether heartbeats between the system and the Quorum Witness are properly sent and received
- **secure_connection**: The state of the security established on the connection. The possible values are:

| Value | Meaning |
|---|---|
| `None` | Security is disabled on the Quorum Witness connectivity. |
| `Verified` | The connection security is verified. |
| `Unauthorized` | The system certificate was rejected by the Quorum Witness. |

- **version**: The Quorum Witness software version
- **id**
- **certificate**
- **db_init**: The timestamp (in mono-time) of the last DB initialization
- **first_event_id**: The ID of the first relevant event
- **last_event_id**: The ID of the last relevant event

**Example:**

```
quorum_witness_list name=q1
```

**Output:**

```
Name     ID                 Address     Port   State        Connection    External Name
------   ----------------   ----------  -----  ----------   ----------    -------------
q1       quorum_witness_id  192.0.2.1   4567   Activated    Up            quorum_witness

Cont:
Db Health
---------
OK
```

| Field ID | Field output | Default position |
|---|---|---|
| **name** | Name | 1 |
| **quorum_id** | ID | 2 |
| **address** | Address | 3 |
| **port** | Port | 4 |
| **state** | State | 5 |
| **connection** | Connection | 6 |
| **external_name** | External Name | 7 |
| **db_health** | Db Health | 8 |
| **heartbeats_ok** | Heartbeating | N/A |
| **use_secure** | Use Secure | N/A |
| **secure_connection** | Secure Connection | N/A |
| **version** | Version | N/A |
| **id** | ID | N/A |
| **certificate** | Certificate | N/A |
| **db_init** | DB Init | N/A |
| **first_event_id** | First Event Id | N/A |
| **last_event_id** | Last Event Id | N/A |

## Access control

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

# Updating a Quorum Witness definition

Use the **quorum_witness_update** command to update a Quorum Witness definition.

```
quorum_witness_update name=qw_name [ certificate=new_qw_certificate ] [ address=new_qw_address ]
 [ port=new_qw_port ]
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **name** | Object name | The Quorum Witness internal name. | Y | N/A |
| **certificate** | N/A | The public certificate or certificate chain of the Quorum Witness (see below for details). | N | Current value. |
| **address** | N/A | The Quorum Witness address: IPv4, IPv6 (full format only) or DNS name. | N | Current value. |
| **port** | Positive integer | A new communication port of the Quorum Witness. | N | Current value. |

As a prerequisite for successfully completing this command, the Quorum Witness must be deactivated.

**The certificate parameter:**

The value of the certificate parameter is the content of a PEM file with asterisks instead of newlines. Chained certificates are supported. The total maximal length of a PEM file holding chained certificates (leaf first, root last) is 15360 characters (including the asterisk characters). In Windows, you can drag-and-drop a PEM file from the Windows Explorer to the appropriate location in the CLI session window; the content will be added automatically.

**Example:**

```
quorum_witness_update name=q1 address=192.0.2.1
```

**Output:**

```
Command completed successfully.
```

## Access control

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

## Return codes

- **QUORUM_WITNESS_BAD_NAME**

The Quorum Witness name does not exist.

- **QUORUM_WITNESS_ADDRESS_ALREADY_EXISTS**

  The Quorum Witness IPv4/6 address or DNS name with same port number already exists.

  **Troubleshooting:** Use the already configured Quorum Witness or change the address or port number.
- **CANT_UPDATE_ACTIVATED_QUORUM_WITNESS**

  Cannot update an activated Quorum Witness.
- **SSL_CERTIFICATE_HAS_EXPIRED**

  The SSL certificate has expired.
- **SSL_CERTIFICATE_ISSUER_NOT_FOUND**

  The SSL certificate issuer was not found in the certificate chain.
- **SSL_CERTIFICATE_INVALID_FORMAT**

  The SSL certificate format is invalid or corrupted.
- **SSL_CERTIFICATE_CHAIN_EMPTY**

  No certificates were found in the input.
- **SSL_CERTIFICATE_VERIFICATION_INTERNAL_ERROR**

  The SSL certificate verification has failed because of an internal system error.
- **SSL_CERTIFICATE_VERIFICATION_FAILED**

  The SSL certificate chain verification failed.
- **SSL_CERTIFICATE_NOT_YET_VALID**

  The SSL certificate is not yet valid.

## Renaming a Quorum Witness

Use the **quorum_witness_rename** command to rename a Quorum Witness.

```
quorum_witness_rename name=qw_name new_name=qw_new_name
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **name** | Object name | The Quorum Witness internal name. | Y |
| **new_name** | Object name | A new internal name of the Quorum Witness. | Y |

**Example:**

```
quorum_witness_rename name=q1 new_name=q2
```

**Output:**

```
Command completed successfully.
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |

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

### Return codes

- **QUORUM_WITNESS_BAD_NAME**

  The Quorum Witness name does not exist.
- **QUORUM_WITNESS_NAME_ALREADY_EXISTS**

  The Quorum Witness name already exists.

## Deleting a Quorum Witness

Use the **quorum_witness_delete** command to delete a Quorum Witness.

```
quorum_witness_delete name=qw_name
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **name** | Object name | The Quorum Witness internal name. | Y |

The command can be issued only for a Quorum Witness that is not in use (attached to any target) and is deactivated or is being deactivated, that is its state is *Deactivated* or *Deactivating*.

**Example:**

```
quorum_witness_delete name=q1 -y
```

**Output:**

```
Command completed successfully.
```

### Access control

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

### Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_DELETE_QUORUM_WITNESS**

Are you sure you want to delete Quorum Witness *Quorum Witness*?

### Return codes

- **QUORUM_WITNESS_BAD_NAME**

  The Quorum Witness name does not exist.
- **CANT_DELETE_AN_ACTIVATED_QUORUM_WITNESS**

  Cannot delete an activated Quorum Witness.

  **Troubleshooting:** Deactivate the Quorum Witness and then try again.
- **CANT_DELETE_QUORUM_WITNESS_IN_USE**

  Cannot delete a Quorum Witness when it is in use by a target.

## Activating a Quorum Witness

Use the **quorum_witness_activate** command to activate a Quorum Witness.

```
quorum_witness_activate name=qw_name
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **name** | Object name | The Quorum Witness internal name. | Y |

This commands activates a Quorum Witness and starts heartbeat and status communication between the system and the Quorum Witness. The command is asynchronous, its completion is tracked by return codes.

**Example:**

```
quorum_witness_activate name=q1
```

**Output:**

```
Command completed successfully.
```

### Access control

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

### Return codes

- **MAX_ACTIVE_QUORUM_WITNESSES_REACHED**

  The maximum number of active Quorum Witness instances is already reached.

**Troubleshooting:** Deactivate an active Quorum Witness and then try again.

- **QUORUM_WITNESS_BAD_NAME**

  The Quorum Witness name does not exist.

- **QUORUM_WITNESS_IS_ALREADY_ACTIVE**

  The Quorum Witness is already active or is being activated.

  **Troubleshooting:** Wait for the activation process to complete. If this issue persists, contact IBM Support.

- **QUORUM_WITNESS_MISSING_SERVICE_CERTIFICATE**

  No valid certificate is defined for the Quorum Witness service.

  **Troubleshooting:** Use PKI commands to define the certificate for the Quorum Witness service.

- **SSL_CERTIFICATE_HAS_EXPIRED**

  The SSL certificate has expired.

# Deactivating a Quorum Witness

Use the **quorum_witness_deactivate** command to deactivate a Quorum Witness.

```
quorum_witness_deactivate name=qw_name
```

## Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **name** | Object name | The Quorum Witness internal name. | Y |

This commands deactivates a Quorum Witness and stops the heartbeat and status communication between the system and the Quorum Witness. It is not allowed to deactivate a Quorum Witness that is in use by an active HyperSwap relation.

**Example:**

```
quorum_witness_deactivate name=q1 -y
```

**Output:**

```
Command completed successfully.
```

## Access control

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

## Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_DEACTIVATE_QUORUM_WITNESS**

Are you sure you want to deactivate Quorum Witness *Quorum Witness*?

### Return codes

- **QUORUM_WITNESS_BAD_NAME**

  The Quorum Witness name does not exist.
- **QUORUM_WITNESS_IS_ALREADY_INACTIVE**

  The Quorum Witness is already inactive or is being deactivated.

  **Troubleshooting:** Wait for the deactivation process to complete. If this issue persists, contact IBM Support.
- **CANT_DEACTIVATE_QUORUM_WITNESS_IN_USE**

  Cannot deactivate a Quorum Witness that is in use by a target with an active HyperSwap relationship.

## Listing the Quorum Witness connection status

Use the **quorum_witness_connections_list** command to list the status of the connections with the Quorum Witness.

```
quorum_witness_connections_list [ name=qw_name ] [ module=module_id ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **name** | Object name | The internal name of the Quorum Witness to be listed. | N | All Quorum Witness connections. |
| **module** | N/A | The ID of the module to be listed. | N | All modules to which Quorum Witness is connected. |

The command output includes the following fields:

- **module_id**. This is the ID of the module on which the Quorum Witness is running.
- **name**. This is the Quorum Witness internal name.
- **connection**. This is state of the connection with the Quorum Witness. The possible values are:

| Value | Meaning |
|---|---|
| Band content encoding | N/A |
| Cannot connect to peer | N/A |
| Cannot resolve the proxy address | N/A |
| Cannot resolve the host address | N/A |
| Connected | A Quorum Node is successfully connected to the Quorum Witness. |
| Connection timeout | The Quorum Witness may be down or the address/port pair is incorrect. |
| Failed to chunk data | N/A |
| Failed to initialize | The HTTP connection initialization failed. |
| HTTP communication error | N/A |
| HTTP Post communication error | N/A |
| Malformed URL | N/A |

| Value | Meaning |
|---|---|
| Out of memory | N/A |
| Quorum node has failed | N/A |
| Secure connection issue | See details in the **secure_connection** field. |
| Too many redirects | N/A |
| Transient error | N/A |

- **secure_connection**: The state of the security established on the connection. The possible values are:

| Value | Meaning |
|---|---|
| None | Security is disabled on the Quorum Witness connectivity. |
| CA certificate is unusable | The provided Quorum Witness certificate is incorrect. |
| Could not load CACERT file, missing or wrong format | N/A |
| Could not load CRL file, missing or wrong format | N/A |
| Failed initializing secure communication | N/A |
| Failed to shut down the SSL/TLS connection | N/A |
| General SSL/TLS failure | N/A |
| Issuer check failed | N/A |
| Peer verification failed | N/A |
| Problem with local certificate | N/A |
| Secure cipher error | N/A |
| Unauthorized | The system certificate was rejected by the Quorum Witness. |
| Verified | The connection security is verified. |

- **heartbeating**: Indicates whether the module heartbeats are successfully processed by the Quorum Witness.

**Example:**

```
quorum_connection_list name=q1 module=1:Module:3
```

**Output:**

```
Module ID    Name    Connection   Secure Connection   Heartbeating
------------ ------- ------------ ------------------- --------------
1:Module:3   q1      Connected    Verified            yes
```

| Field ID | Field output | Default position |
|---|---|---|
| **module_id** | Module ID | 1 |
| **name** | Name | 2 |
| **connection** | Connection | 3 |
| **secure_connection** | Secure Connection | 4 |
| **heartbeats_ok** | Heartbeating | 5 |

### Access control

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

## Getting Quorum Witness information

Use the **quorum_witness_info_get** command to get Quorum Witness information from a Quorum Witness.

```
quorum_witness_info_get name=qw_name
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **name** | Object name | The Quorum Witness internal name. | Y |

The command output includes the following fields:

- **internal_name**: The Quorum Witness internal name
- **quorum_id**: A globally unique Quorum Witness ID
- **name**: The Quorum Witness external name
- **version**: The Quorum Witness software version

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

- **db_state**: JSON string (up to 256 bytes), representing the DB disk usage
- **host_info**: The type and the version of the host's operating system, on which the Quorum Witness is running
- **protocol**: The protocol version
- **network_load**: JSON string (up to 128 bytes), representing the count of connections and requests
- **cpu_load**: JSON string (up to 128 bytes), representing the CPU consumption percentage
- **last_event_id**: The ID of the last relevant event
- **db_init**: The timestamp (in mono-time) of the last DB initialization

**Example:**

```
quorum_witness_info_get name=q1
```

**Output:**

```
Name   ID                                External Name   Version
------ --------------------------------- --------------- ---------
q1     7f2f35834fea48b6b335aa09fd9179ae  FVT-QW104       1.0.0
```

| Field ID | Field output | Default position |
|---|---|---|
| **internal_name** | Name | 1 |
| **quorum_id** | ID | 2 |
| **name** | External Name | 3 |
| **version** | Version | 4 |
| **db_state** | DB state | N/A |
| **host_info** | Host | N/A |
| **protocol** | Protocol Version | N/A |
| **network_load** | Network Load | N/A |
| **cpu_load** | CPU Load | N/A |
| **last_event_id** | Last Event ID | N/A |
| **db_init** | DB Initialization Timestamp | N/A |

## Access control

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

## Return codes

- **QUORUM_WITNESS_BAD_NAME**

  The Quorum Witness name does not exist.
- **QUORUM_WITNESS_RESPONSE_TIMEOUT**

  No response from Quorum Witness *Quorum Witness* within the designated timeout period.

# 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**

  The volume name does not exist.
- **VOLUME_NO_DM**

  The local volume does not have data migration definitions.
- **TARGET_NOT_CONNECTED**

  There is currently no connection to the target system.
- **REMOTE_VOLUME_NO_LUN**

  The remote volume's LUN is unavailable.
- **REMOTE_VOLUME_NO_READ_ACCESS**

  The remote volume cannot be read.
- **REMOTE_VOLUME_NO_WRITE_ACCESS**

  The remote volume is write-protected.

© Copyright IBM Corp. 2016, 2022

- **BAD_REMOTE_VOLUME_SIZE**

  The primary and secondary volumes contain a different number of blocks.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support

# Deactivating data migration

Use the **dm_deactivate** command to deactivate the data migration process.

```
dm_deactivate vol=VolName
```

## Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **vol** | Object name | The local volume on which the data migration process is to be deactivated. | Y |

Hosts are not served while the data migration process is inactive.

This command has no effect if the data migration process is already inactive.

## Access control

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

## Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_DEACTIVATE_DATA_MIGRATION**

  Deactivation will stop all applications. After deactivation, data migration can be deleted.
- **ARE_YOU_SURE_YOU_WANT_TO_DEACTIVATE_SOURCE_UPDATING_DATA_MIGRATION**

  Deactivation may cause loss of access to hosts, and will stop all applications. After deactivation, data migration can be deleted.

## Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **VOLUME_NO_DM**

  The local volume does not have data migration definitions.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

**Troubleshooting:** Contact IBM Support

# Defining data migration configuration

Use the **dm_define** command to define a data migration configuration.

```
dm_define vol=VolName target=TargetName lun=SourceLUN source_updating=<yes|no> [ create_vol=<yes|
no> ]
 [ pool=PoolName ]
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **vol** | Object name | Data migration destination volume on the local system. | Y | N/A |
| **target** | Object name | Remote system containing the source volume. | Y | N/A |
| **lun** | Integer | LUN of the source volume. | Y | N/A |
| **source_updating** | Boolean | Specifies whether to use source volume updating. | Y | N/A |
| **create_vol** | Boolean | A Boolean that determines whether to create a new volume or to use an existing one. | N | No |
| **pool** | Object name | Name of the storage pool to contain the volume. Used only when creating a volume. Mandatory when creating a volume. | N | N/A |

This command defines a data migration relation between a local volume and a remote volume. According to this definition, the local volume should reflect the remote volume.

After this configuration has been defined, it can be tested using the **dm_test** command (see Testing the data migration definition) and then activated using the **dm_activate** command (see Activating data migration). After this activation, hosts can read and write to this volume, and these operations are reflected on the remote volume.

The remote volume may be inaccessible when the command is executed. In this case, the definition is only used when data migration is tested.

The local system acts as a host to the remote system. The remote system should be configured to make the remote volume accessible to the local system through the specified LUN.

If **source updating** is specified, each write to the local volume is reflected as a write to the remote volume. Otherwise, writes on the local volume are not reflected and the remote volume is not changed.

The local volume must be formatted.

If **create_vol** is set to *yes*, the volume is created. In this case the size of the newly created volume is identical to the size of the source volume. When creating a volume, a pool name must be specified. Creating a volume fails if there is no connectivity to the target since the volume's size is unknown.

## Access control

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

## Warnings

- **VOLUME_SIZE_VERY_LARGE_ARE_YOU_SURE**

  The volume size is very large. It may not be possible to mirror this volume to older versions of the storage system. Are you sure?

## Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **TARGET_BAD_NAME**

  The target name does not exist.
- **VOLUME_IS_SNAPSHOT**

  The operation is not permitted on snapshots.
- **VOLUME_HAS_MIRROR**

  A mirror is defined for this volume.
- **VOLUME_BELONGS_TO_CG**

  The volume belongs to a consistency group.
- **VOLUME_HAS_DATA_MIGRATION**

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

  The volume has snapshots.
- **VOLUME_NOT_FORMATTED**

  The local volume is not formatted.
- **VOLUME_EXISTS**

  The volume name already exists.
- **POOL_DOES_NOT_EXIST**

  The storage pool does not exist.
- **VOLUME_BAD_PREFIX**

  The volume name has a reserved prefix.
- **NOT_ENOUGH_SPACE**

  No space to allocate for the volume's current usage.
- **MAX_VOLUMES_REACHED**

  The maximum allowed number of volumes is already reached.

- **REMOTE_VOLUME_NO_LUN**

  The remote volume's LUN is unavailable.
- **TARGET_NOT_CONNECTED**

  There is currently no connection to the target system.
- **VOLUME_CANNOT_HAVE_ZERO_SIZE**

  The volume size cannot be zero.
- **ILLEGAL_LUN**

  The LUN is out of range.
- **TARGET_IS_MIRRORING**

  The target machine is defined for remote mirroring only.
- **NO_ONLINE_MIGRATION_WITHOUT_SOURCE_UPDATING**

  Data migration without automatic migration must be defined as a source update.
- **MIGRATION_ALREADY_DEFINED_FOR_LUN**

  Data migration is already defined from lun *LUN* of target '*Target*'.
- **VOLUME_SIZE_ABOVE_LIMIT**

  The specified volume size is above the limit.
- **INVALID_SLICE_OFFSET**

  Slice offset is illegal.
- **ENCRYPTION_IN_PROGRESS**

  The system is in the process of changing the encryption activation state.
- **DOMAIN_MAX_VOLUMES_REACHED**

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

  The maximum number of remote volumes (mirror/migration) is already reached.

  **Troubleshooting:** Delete unnecessary data migration objects.
- **DOMAIN_MAX_DMS_REACHED**

  The domain exceeds the maximum allowed number of data migrations.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support
- **VOLUME_TOO_BIG**

  No space to allocate to the volume.

## Deleting a data migration process

Use the **dm_delete** command to delete the data migration process.

```
dm_delete vol=VolName [ force_delete=<yes|no> ]
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **vol** | Object name | The name of the volume whose data migration process is to be deleted. | Y | N/A |
| **force_delete** | Boolean | When set to *yes*, forces the deletion even if data migration is not complete. | N | no |

This command deletes a data migration process and its configuration after it is completed and the target has a full copy of the data. After deleting a data migration process, the local volume must be explicitly deleted manually.

**Using the force_delete parameter**

The **force_delete** parameter enables deletion of a data migration process when only the source has a complete copy of data. This is mostly necessary if the data migration configuration is incorrect: a wrong source volume is selected or the defined volume (block) size is unsuitable.

When data migration is source updating and, therefore, the source has full data, the **force_delete** parameter is allowed as well.

It is currently impossible to delete a data migration process when neither the source nor the target have complete data.

## Access control

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

## Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **VOLUME_NO_DM**

  The local volume does not have data migration definitions.
- **DM_IS_NOT_SYNCHRONIZED**

  The data migration process has not been completed.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support
- **DM_OPERATION_NOT_ALLOWED_NEITHER_SOURCE_NOR_TARGET_CONTAIN_FULL_DATA**

  Data migration was set up without source update and synchronization has not completed yet. As a result, neither the source volume nor the target volume contains a full image of the data.

# Listing data migration statuses

Use the **dm_list** command to list data migration configuration and status.

```
dm_list [ vol=VolName ] [ domain=DomainName ]
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **vol** | Object name | The name of the volume to be listed. | N | All data migration volumes. |
| **domain** | Object name | The domain name. | N | All Domains |

This command lists all data migration configuration and statuses, including the following information:

- Volume name
- Target name
- LUN
- Volume size (GB)
- Migration completed (GB)
- Migration activation (active/inactive)
- Migration status (synchronized, unsynchronized)
- Migration remaining (GB)
- Migration remaining (%)
- Estimated time to completion

| Field ID | Field output | Default position |
|---|---|---|
| **local_volume_name** | Local Volume | 1 |
| **target_name** | Remote System | 2 |
| **remote_volume_lun** | Remote LUN | 3 |
| **active** | Active | 4 |
| **sync_state** | Status | 5 |
| **connected** | Target Connected | 6 |
| **size_to_synchronize** | Size To Sync (MiB) | N/A |
| **operational** | Operational | N/A |
| **sync_progress** | Sync Progress (%) | N/A |
| **start_migration_automatically** | Start Data Migration Automatically | N/A |
| **arch** | Remote Arch | N/A |

## Access control

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

| User Category | Permission |
|---|---|
| Technicians | Disallowed |

# Testing the data migration definition

Use the **dm_test** command to test the data migration configuration.

```
dm_test vol=VolName
```

## Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **vol** | Object name | Destination volume for data migration testing. | Y |

Command return codes indicate the types of test failures that may occur. Once a test is successful, then data migration can be activated.

If source updating is not defined for this data migration, the writing is not tested.

## Access control

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

## Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **VOLUME_NO_DM**

  The local volume does not have data migration definitions.
- **TARGET_NOT_CONNECTED**

  There is currently no connection to the target system.
- **REMOTE_VOLUME_NO_LUN**

  The remote volume's LUN is unavailable.
- **REMOTE_VOLUME_NO_READ_ACCESS**

  The remote volume cannot be read.
- **REMOTE_VOLUME_NO_WRITE_ACCESS**

  The remote volume is write-protected.
- **BAD_REMOTE_VOLUME_SIZE**

  The primary and secondary volumes contain a different number of blocks.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

**Troubleshooting:** Contact IBM Support

# 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**

  The volume size is very large. It may not be possible to mirror this volume to older versions of the storage system. Are you sure?

### Return codes

- **VOLUME_NOT_APPLICABLE_FOR_OLVM**

  The volume is not applicable to IBM Hyper-Scale Mobility.
- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **VOLUME_IS_SNAPSHOT**

  The operation is not permitted on snapshots.
- **TARGET_BAD_NAME**

© Copyright IBM Corp. 2016, 2022

The target name does not exist.

- **VOLUME_IS_MASTER**

  This local volume is already defined as a primary volume.
- **TARGET_BAD_TYPE**

  The target machine is not an XIV machine.
- **TARGET_NO_ACCESS**

  No access permissions to the secondary machine.
- **TARGET_NOT_CONNECTED**

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

  There is currently no connection from the target system.
- **MAX_MIGRATIONS_REACHED**

  The maximum number of migrations is already reached.
- **REMOTE_MAX_MIGRATIONS_REACHED**

  The maximum number of migrations is already reached on the remote machine.
- **REMOTE_POOL_DOES_NOT_EXIST**

  The pool does not exist on the remote machine.
- **BAD_REMOTE_VOLUME_SIZE**

  The primary and secondary volumes contain a different number of blocks.
- **NOT_ENOUGH_SPACE_ON_REMOTE_MACHINE**

  Not enough free space to set the requested size of the secondary volume.
- **REMOTE_VOLUME_EXISTS**

  The secondary volume with the indicated name already exists. The name cannot be reused.
- **REMOTE_VOLUME_IS_MASTER**

  A volume on the remote machine is already defined as primary.
- **REMOTE_VOLUME_IS_SLAVE**

  A volume on the remote machine is already defined as secondary.
- **REMOTE_MAX_VOLUMES_REACHED**

  The maximum number of volumes on the remote machine is already reached.
- **TIMEOUT**

  A remote operation was not completed in time.
- **VOLUME_BAD_PREFIX**

  The volume name has a reserved prefix.
- **REMOTE_VOLUME_HAS_DATA_MIGRATION**

  Data migration is already defined for the secondary volume.
- **VOLUME_HAS_OLVM**

  An IBM Hyper-Scale Mobility relationship is defined for this volume.
- **VOLUME_HAS_HA**

  This operation is forbidden on a volume with a HyperSwap relation.
- **TARGET_VOLUME_HAS_OLVM**

  This target volume is part of an IBM Hyper-Scale Mobility relationship.

- **TARGET_VOLUME_HAS_HA**

  This operation is forbidden, if the target volume is a peer in a HyperSwap relation.
- **REMOTE_VOLUME_LOCKED**

  The secondary volume is locked.
- **VOLUME_HAS_MIRRORING_SNAPSHOTS**

  The volume has snapshots created by a previous mirroring process.
- **REMOTE_MAX_MIRROR_CAPACITY_REACHED**

  The maximum capacity for mirrored volumes is already reached on the remote machine.
- **TARGET_DOES_NOT_ACCEPT_XIV_COMMANDS**

  The target system does not accept XIV management commands.
- **MAX_VOLUMES_REACHED**

  The maximum allowed number of volumes is already reached.
- **VOLUME_LOCKED**

  The volume is locked.
- **NO_ASYNC_IN_THIN_PROVISIONED_POOL**

  A thin-provisioned pool cannot contain volumes with asynchronous mirroring.
- **BAD_REMOTE_VOLUME_NAME**

  The secondary volume name does not exist.
- **REMOTE_VOLUME_SIZE_ABOVE_LIMIT**

  The specified volume size is above the limit of the remote machine.
- **MIRROR_RETRY_OPERATION**

  There is an operation in progress on this mirror.

  **Troubleshooting:** Retry the command in a few seconds.
- **ELECTRONIC_LICENSE_NOT_APPROVED**

  Operation blocked until Electronic license approval

  **Troubleshooting:** Please retrieve Electronic license version and accept it
- **VOLUME_NOT_FORMATTED**

  The local volume is not formatted.
- **MIRROR_TYPE_INCOMPATIBLE_WITH_TARGET**

  A mirror of this type is not supported between the system versions of the specified peers.
- **OLVM_INCOMPATIBLE_TARGET_VERSION**

  IBM Hyper-Scale Mobility is not supported between the system versions of the specified peers.
- **VOLUME_TOO_BIG**

  No space to allocate to the volume.
- **VOLUME_SIZE_ABOVE_LIMIT**

  The specified volume size is above the limit.
- **INVALID_SLICE_OFFSET**

  Slice offset is illegal.
- **VOLUME_IS_OLVM_PROXY**

  The volume is in an IBM Hyper-Scale Mobility Proxy phase.
- **VOLUME_IS_SLAVE**

The volume is defined as a secondary volume.

- **REMOTE_VOLUME_IS_SNAPSHOT**

  The secondary volume is a snapshot.

- **VOLUME_EXISTS**

  The volume name already exists.

- **SLAVE_VOLUME_NOT_FORMATTED**

  The secondary volume is not formatted.

- **VOLUME_BELONGS_TO_CG**

  The volume belongs to a consistency group.

- **VOLUME_HAS_DATA_MIGRATION**

  Data Migration is defined for this volume.

- **MAX_MIRRORS_REACHED**

  The maximum number of mirrors is already reached.

- **VOLUME_CANNOT_HAVE_ZERO_SIZE**

  The volume size cannot be zero.

- **ASYNC_MIRROR_REMOTE_RPO_TOO_SHORT**

  The specified remote RPO is too short.

- **POOL_DOES_NOT_EXIST**

  The storage pool does not exist.

- **REMOTE_VOLUME_NOT_APPLICABLE_FOR_OLVM**

  The remote volume is not applicable to IBM Hyper-Scale Mobility.

- **REMOTE_SCHEDULE_DOES_NOT_EXIST**

  The specified schedule does not exist on the remote machine.

- **OLVM_DOES_NOT_SUPPORT_ISCSI_TARGET**

  IBM Hyper-Scale Mobility does not support ISCSI targets.

- **ASYNC_MIRROR_REMOTE_RPO_TOO_LONG**

  The specified remote RPO is too long.

- **OPERATION_DENIED_OBJECT_MANAGED**

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

- **ENCRYPTION_IN_PROGRESS**

  The system is in the process of changing the encryption activation state.

- **MAX_OLVM_REACHED**

  The maximum allowed number of IBM Hyper-Scale Mobility relationships is already reached.

- **DOMAIN_MAX_MIRRORS_REACHED**

  The domain exceeds the maximum allowed number of mirrors.

- **REMOTE_DOMAIN_MAX_MIGRATIONS_REACHED**

  The maximum number of migrations is already reached in the remote machine domain.

- **DOMAIN_HAS_NO_ACCESS_TO_TARGET**

  The domain has no access to the target.

- **REMOTE_DOMAIN_HAS_NO_ACCESS_TO_TARGET**

The secondary machine domain has no access to the target.

- **DOMAIN_MAX_VOLUMES_REACHED**

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

  The maximum number of volumes in the remote machine domain is already reached.
- **MULTISITE_MAX_NUM_OF_MIRRORS_REACHED**

  Failed to create the relation, the maximum allowed number of relations is already exceeded.
- **REMOTE_VOLUME_TWO_SYNC_MIRRORS_NOT_ALLOWED**

  Two synchronous mirrors were detected on the remote volume. This is not allowed.
- **REMOTE_VOLUME_IS_MIRROR_MASTER**

  The volume is primary in a mirror relationship, and cannot be secondary!
- **REMOTE_VOLUME_MIRROR_LOOP_DETECTED**

  A mirror loop was detected on the remote volume. This means that there is a mirror on the remote system, whose target is this system. Therefore, you cannot create a mirror with this target on this system.
- **REMOTE_VOLUME_MASTER_ASYNC_MIRROR_DETECTED**

  An asynchronous primary mirror was detected on the remote volume. The operation not allowed.
- **MULTISITE_INCOMPATIBLE_TARGET_VERSION**

  Multi-site is not supported between the system versions of the specified peers.
- **REMOTE_VOLUME_HAS_MIRRORING_SNAPSHOTS**

  The remote volume has snapshots created by a previous mirroring process.
- **VOLUME_HAS_MULTIPLE_RELATIONS**

  The volume has multiple relations. The operation is not allowed or a target must be specified.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support
- **REMOTE_DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier of the remote system is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support
- **SYSTEM_OUT_OF_PHYSICAL_SPACE**

  The operation not allowed while the system is out of physical space.
- **REMOTE_SYSTEM_OUT_OF_PHYSICAL_SPACE**

  The operation not allowed while the remote system is out of physical space.
- **TARGET_CONFIGURATION_AND_CODE_VERSION_DO_NOT_SUPPORT_OPERATION**

  The target configuration and code version do not support this operation.

## Activating a volume migration

Use the **olvm_activate** command to activate an IBM Hyper-Scale Mobility migration for a defined an IBM Hyper-Scale Mobility relation.

```
olvm_activate vol=VolName
```

## Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **vol** | Object name | IBM Hyper-Scale Mobility source volume. | Y |

This command is issued on the source.

## Access control

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

## Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **COMMAND_NOT_SUPPORTED_FOR_OLVM_SOURCE_IN_THIS_STATE**

  The source is in an unsupported IBM Hyper-Scale Mobility state.
- **DOMAIN_MAX_VOLUMES_REACHED**

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

  The mirror's local configuration does not match its remote configuration.
- **MIRROR_CONS_GROUP_MEMBERSHIP_MISMATCH**

  The mirrored consistency group contains different volumes on the primary and secondary machines. This problem occurs whenever the cg_add_vol or cg_remove_vol commands were previously issued, and the primary machine did not receive an acknowledgment from the secondary machine until the command timed out, or due to any other unexpected failure.
- **MIRROR_RETRY_OPERATION**

  There is an operation in progress on this mirror.

  **Troubleshooting:** Retry the command in a few seconds.
- **COMMAND_NOT_SUPPORTED_FOR_OLVM_DESTINATION_IN_THIS_STATE**

  The destination is in an unsupported IBM Hyper-Scale Mobility state.
- **MIRROR_SIZE_MISMATCH**

  The secondary and primary volume sizes are different.
- **REMOTE_VOLUME_IS_MASTER**

  A volume on the remote machine is already defined as primary.
- **REMOTE_MAX_VOLUMES_REACHED**

  The maximum number of volumes on the remote machine is already reached.
- **VOLUME_NOT_DEFINED_FOR_OLVM**

  The volume does not have IBM Hyper-Scale Mobility definitions.

- **OLVM_ALREADY_ACTIVE**

  The IBM Hyper-Scale Mobility relationship is already active.
- **MAX_VOLUMES_REACHED**

  The maximum allowed number of volumes is already reached.
- **REMOTE_MIRROR_IS_STANDBY**

  The remote mirror is marked as Standby.
- **VOLUME_HAS_MULTIPLE_RELATIONS**

  The volume has multiple relations. The operation is not allowed or a target must be specified.
- **REMOTE_DOMAIN_MAX_VOLUMES_REACHED**

  The maximum number of volumes in the remote machine domain is already reached.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support
- **REMOTE_DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier of the remote system is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support
- **REMOTE_SYSTEM_OUT_OF_PHYSICAL_SPACE**

  The operation not allowed while the remote system is out of physical space.
- **OLVM_INCOMPATIBLE_TARGET_VERSION**

  IBM Hyper-Scale Mobility is not supported between the system versions of the specified peers.

# Deactivating IBM Hyper-Scale Mobility migration

Use the **olvm_deactivate** command to deactivate IBM Hyper-Scale Mobility migration for a defined IBM Hyper-Scale Mobility relation.

```
olvm_deactivate vol=VolName
```

## Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **vol** | Object name | The source volume. | Y |

This command is issued on the source.

## Access control

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

### Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_DEACTIVATE_OLVM**

  Are you sure you want to deactivate IBM Hyper-Scale Mobility?

### Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **REMOTE_VOLUME_IS_MASTER**

  A volume on the remote machine is already defined as primary.
- **COMMAND_NOT_SUPPORTED_FOR_OLVM_SOURCE_IN_THIS_STATE**

  The source is in an unsupported IBM Hyper-Scale Mobility state.
- **VOLUME_NOT_DEFINED_FOR_OLVM**

  The volume does not have IBM Hyper-Scale Mobility definitions.
- **COMMAND_NOT_SUPPORTED_FOR_OLVM_DESTINATION_IN_THIS_STATE**

  The destination is in an unsupported IBM Hyper-Scale Mobility state.
- **OLVM_ALREADY_INACTIVE**

  The IBM Hyper-Scale Mobility relationship is already inactive.
- **REMOTE_MIRROR_IS_STANDBY**

  The remote mirror is marked as Standby.
- **VOLUME_HAS_MULTIPLE_RELATIONS**

  The volume has multiple relations. The operation is not allowed or a target must be specified.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support

## Aborting a defined or activated IBM Hyper-Scale Mobility process

Use the **olvm_abort** command to abort a defined or activated IBM Hyper-Scale Mobility process.

```
olvm_abort < vol=VolName [ force_abort=<yes|no> | force_abort_on_destination=<yes|no> ] >
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **vol** | Object name | The source volume. | Y | N/A |
| **force_abort** | Boolean | Determines whether to delete an IBM Hyper-Scale Mobility relation on the source. | N | No |
| **force_abort_on_destination** | Boolean | Determine whether to delete an IBM Hyper-Scale Mobility relation on the destination. | N | No |

This command is issued on the source and has the option to abort the IBM Hyper-Scale Mobility process either from the source or from the destination.

Once issued, the source and destination are no longer part of an IBM Hyper-Scale Mobility relation. IBM Hyper-Scale Mobility attributes are deleted.

### Access control

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

### Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_ABORT_OLVM_RELATIONSHIP_IN_THIS_PHASE**

  Are you sure you want to abort the IBM Hyper-Scale Mobility relationship?
- **ARE_YOU_SURE_YOU_WANT_TO_FORCE_ABORT_OLVM_RELATIONSHIP_IN_THIS_PHASE**

  Are you sure you want to force abort the IBM Hyper-Scale Mobility relationship?

### Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **VOLUME_NOT_DEFINED_FOR_OLVM**

  The volume does not have IBM Hyper-Scale Mobility definitions.
- **COMMAND_NOT_SUPPORTED_FOR_OLVM_SOURCE_IN_THIS_STATE**

  The source is in an unsupported IBM Hyper-Scale Mobility state.
- **COMMAND_NOT_SUPPORTED_FOR_OLVM_DESTINATION_IN_THIS_STATE**

  The destination is in an unsupported IBM Hyper-Scale Mobility state.
- **OLVM_IS_ACTIVE**

  The IBM Hyper-Scale Mobility relationship is active.
- **FORCE_ABORT_NOT_ALLOWED**

  A forced IBM Hyper-Scale Mobility abort is not allowed.
- **VOLUME_HAS_MULTIPLE_RELATIONS**

  The volume has multiple relations. The operation is not allowed or a target must be specified.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support

## Moving the IBM Hyper-Scale Mobility source volume to a Proxy state

Use the **olvm_proxy** command to move the IBM Hyper-Scale Mobility source volume to a Proxy state.

```
olvm_proxy vol=VolName
```

## Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **vol** | Object name | The source volume. | Y |

This command moves the IBM Hyper-Scale Mobility source volume to a Proxy state where the source acts as a proxy to the destination.

The source becomes a proxy and the destination becomes the data 'owner'. Host writes are no longer written to the source and the volume data on the source is freed. The source volume and snapshot data are deleted.

This command is issued on the source.

## Access control

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

## Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_OLVM_PROXY**

  Are you sure you want to move the volume *Volume*to the Proxy state? The source volume and all volume snapshots will be deleted.

## Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **COMMAND_NOT_SUPPORTED_FOR_OLVM_SOURCE_IN_THIS_STATE**

  The source is in an unsupported IBM Hyper-Scale Mobility state.
- **COMMAND_NOT_SUPPORTED_FOR_OLVM_DESTINATION_IN_THIS_STATE**

  The destination is in an unsupported IBM Hyper-Scale Mobility state.
- **VOLUME_NOT_DEFINED_FOR_OLVM**

  The volume does not have IBM Hyper-Scale Mobility definitions.
- **OLVM_LINK_IS_NOT_UP**

  The IBM Hyper-Scale Mobility link is not up. The mapping list cannot be updated.
- **OLVM_PROXY_MOVE_INITIATED**

  IBM Hyper-Scale Mobility volume move to the Proxy state has started.
- **HOST_BAD_NAME**

  The host name does not exist.
- **ISCSI_HOST_ILLEGAL_PORT_NAME**

  The port name for iSCSI Host is illegal.

  **Troubleshooting:** Port names for iSCSI Hosts must contain only printable characters.

- **MAX_PORTS_REACHED**

  The maximum number of ports defined in the system is already reached.
- **HOST_PORT_EXISTS**

  A host with this port ID is already defined.
- **REMOTE_MAX_VIRTUAL_HOSTS_REACHED**

  The maximum number of defined remote virtual hosts is already reached.
- **OLVM_RETRY_OPERATION**

  There is an operation in progress on this OLVM.

  **Troubleshooting:** Retry the command in a few seconds.
- **VOLUME_HAS_MULTIPLE_RELATIONS**

  The volume has multiple relations. The operation is not allowed or a target must be specified.
- **MAX_METADATA_OBJECTS_REACHED**

  The maximum number of metadata objects has been reached.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support
- **LOCAL_VOLUME_HAS_TOO_MANY_METADATA_OBJECTS**

  Local volume has too many metadata objects.

  **Troubleshooting:** Contact IBM Support.

# Deleting an IBM Hyper-Scale Mobility relation

Use the **olvm_delete** command to delete an IBM Hyper-Scale Mobility relation and attributes.

```
olvm_delete vol=VolName [ force_delete=<yes|no> ]
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **vol** | Object name | The volume for IBM Hyper-Scale Mobility abort. | Y | N/A |
| **force_delete** | Boolean | Determines whether to delete an IBM Hyper-Scale Mobility relation on the destination. | N | No |

This command is issued on the source. If there is no communication to the destination, the command can force delete the IBM Hyper-Scale Mobility relation.

## Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Allowed |
| Application administrator | Disallowed |
| Security administrator | Disallowed |

| User Category | Permission |
|---|---|
| Read-only users | Disallowed |
| Technicians | Disallowed |

### Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_DELETE_OLVM_RELATIONSHIP_IN_THIS_PHASE**

  Are you sure you want to delete the IBM Hyper-Scale Mobility relationship?
- **ARE_YOU_SURE_YOU_WANT_TO_FORCE_DELETE_OLVM_RELATIONSHIP_IN_THIS_PHASE**

  Are you sure you want to force delete the IBM Hyper-Scale Mobility relationship?

### Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **VOLUME_NOT_DEFINED_FOR_OLVM**

  The volume does not have IBM Hyper-Scale Mobility definitions.
- **COMMAND_NOT_SUPPORTED_FOR_OLVM_SOURCE_IN_THIS_STATE**

  The source is in an unsupported IBM Hyper-Scale Mobility state.
- **COMMAND_NOT_SUPPORTED_FOR_OLVM_DESTINATION_IN_THIS_STATE**

  The destination is in an unsupported IBM Hyper-Scale Mobility state.
- **FORCE_DELETE_NOT_ALLOWED**

  A forced deletion of the IBM Hyper-Scale Mobility relation is not allowed.
- **VOLUME_IS_MAPPED**

  The volume mapped to a host cannot be deleted.
- **VOLUME_HAS_MULTIPLE_RELATIONS**

  The volume has multiple relations. The operation is not allowed or a target must be specified.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support

## Listing the IBM Hyper-Scale Mobility status

Use the **olvm_list** command to list the IBM Hyper-Scale Mobility configuration and status.

```
olvm_list [ vol=VolName ] [ domain=DomainName ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **vol** | Object name | The volume name to be listed. | N | Displays details for IBM Hyper-Scale Mobility relations in the local system. |
| **domain** | Object name | The domain name. | N | All Domains |

This command is issued on the source. The output includes the following information:

- Volume name
- Role (Source, Destination)
- Remote System
- Active (Yes, No)
- Phase (Migration, Proxy-Ready, Proxy)
- State
- Link Up

| Field ID | Field output | Description | Default position |
|---|---|---|---|
| **name** | Volume name | N/A | 1 |
| **role** | Role | N/A | 2 |
| **target_name** | Remote System | N/A | 3 |
| **active** | Active | N/A | 4 |
| **phase** | Phase | N/A | 5 |
| **state** | State | N/A | 6 |
| **connected** | Link Up | N/A | 7 |
| **sync_progress** | Sync Progress (%) | N/A | N/A |
| **size_to_synchronize** | Size To Sync (MiB) | N/A | N/A |
| **estimated_sync_time** | Est. remaining time (sec) | N/A | N/A |
| **mirror_error** | Mirror Error | No Error, Secondary pool exhausted, Configuration error or No thin provisioning resources | N/A |
| **arch** | Remote Arch | N/A | N/A |

## Access control

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

# 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> ]
 [ internal=<yes|no> ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **description** | String | Description of the event. | Y | N/A |
| **severity** | N/A | Severity of the event. | N | Informational |
| **internal** | Boolean | Defines whether this is an XIV internal custom event. | N | no |

This command can be used to either generate an event from a user application or host side software, or to test the event notification procedures.

**Example:**

```
custom_event description="Test started"
```

**Output:**

```
Command completed successfully.
```

### Access control

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

## Generating a CSS product event

Use the **css_product_event** command to generate a CSS (Cloud Storage Solutions) custom event.

```
css_product_event product=Product version=Version server=Server platform=Platform action=Action
properties=Properties [ severity=<INFORMATIONAL|WARNING|MINOR|MAJOR|CRITICAL> ]
```

© Copyright IBM Corp. 2016, 2022

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **product** | String | Product name. | Y | N/A |
| **version** | String | Version information. | Y | N/A |
| **server** | String | Server name. | Y | N/A |
| **platform** | String | Platform information. | Y | N/A |
| **action** | String | Action information. | Y | N/A |
| **properties** | String | Properties information. | Y | N/A |
| **severity** | N/A | Severity of the event. | N | Informational |

This command can be used to either generate an event from a user application or host side software, or to test the event notification procedures.

**Example:**

```
css_product_event product=product_name version=version_info server=server_info
platform=platform_name action=action_name properties=properties_details
```

**Output:**

```
Command completed successfully.
```

### Access control

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

## Defining a new event notification destination

Use the **dest_define** command to define a new destination for event notifications.

```
dest_define dest=DestName type=<SNMP|EMAIL|SMS|HTTPS> < snmp_manager=SNMPManager | <
uri=HTTPSaddress [ proxy=
ProxyAddress [ proxy_port=ProxyPortNum ] ]
 > | email_address=email | <area_code=AreaCode    number=PhoneNumber> |    user=UserName> [
smtpgws=<SMTPGW1[,SMTPGW2]...|ALL> | smsgws=<SMSGW1[,SMSGW2]...|ALL> ]
 [ heartbeat_test_hour=HH:MM [ heartbeat_test_days=Day ] ] [ domain=DomainList ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **dest** | Object name | Destination name. | Y | N/A |
| **type** | Enumeration | Destination type for event notifications: be email, SMS, HTTPS or SNMP. | Y | N/A |

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **snmp_manager** | N/A | IP address or DNS name of the SNMP manager. | N | N/A |
| **uri** | N/A | IP address or DNS name of the HTTPS server. If a port different from the default should be used, specify it here. | N | N/A |
| **proxy** | N/A | IP address or DNS name of the proxy server to send HTTPS over. | N | None |
| **proxy_port** | Integer | Proxy port number to send HTTPS through. The default is 1080. | N | None |
| **email_address** | N/A | Email address. | N | N/A |
| **smtpgws** | Object name | List of SMTP gateways to be used. | N | ALL (all gateways). |
| **area_code** | N/A | Area code of the cellular number for SMS notification. Use digits, '-' or '.' | N | N/A |
| **number** | N/A | Cellular number for SMS notification. Use digits, '-' or '.' | N | N/A |
| **smsgws** | Object name | SMS gateways to be used for this destination. | N | ALL (all gateways). |
| **user** | Object name | User name, where the user's email or phone are used. | N | N/A |
| **heartbeat_test_hour** | N/A | The hour for periodic heartbeat testing in the format HH:MM | N | No heartbeat |
| **heartbeat_test_days** | N/A | List of days for heartbeat testing: a comma-separated list of 3-letter day names (such as "mon", "mon,fri", etc.). | N | No heartbeat |
| **domain** | N/A | Attach the destination to the specified domains. To define more than one domain, separated them with a comma. To specify all existing domains, use "*". | N | none |

This command defines a destination for event notifications. There are four types of destinations: email, SMS, HTTPS and SNMP.

**Note:**

Defining the HTTPS destination can only be done upon an RPQ request.

- *Email* destinations are used for sending notifications via email. When defining a new destination of type Email, either the email address of the recipient must be specified in *email_address* or the user name must be specified in *user* (in this case the email address of that user is used).
- *SMS* destinations are used for sending notifications via SMS to cellular phones. When defining a new destination of type SMS, either the cellular phone number of the destination must be specified in

*number* or the user name must be specified in *user* (in this case the cellular phone number of that user is used). To allow correct formatting, the area code must be separated from the local number.

- *SNMP* destinations are used for sending notifications by SNMP traps to SNMP managers. When defining a new destination of type SNMP, the IP address of the SNMP manager must be specified.
- *HTTPS* destinations are used for sending notifications to HTTPS servers. When defining a new destination of type HTTPS, the IP address of the HTTPS server must be specified.

By default, when sending an email notification, all SMTP gateways specified in the smtpgw_prioritize command (see Prioritizing SMTP gateways) are used, according to the order specified in that command. It is possible to define that sending emails to a specific destination will use specific SMTP gateway or gateways. This is done by specifying the **smtpgws** parameter.

The same logic applies to sending SMS messages. By default, SMS gateways specified in the smsgw_prioritize command (see Prioritizing SMS gateways) are used, according to the order specified in this command. It is possible to define that messages to a specific SMS destination will be sent through a specific SMS gateway or gateways.

**Example:**

```
dest_define dest=adminemail type=EMAIL email_address=storageadmin@example.com
```

**Output:**

```
Command completed successfully.
```

**Example:**

```
dest_define dest=monitoringserver type=SNMP snmp_manager=192.0.2.111
```

**Output:**

```
Command completed successfully.
```

## Access control

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

## Return codes

- **DEST_MAX_REACHED**

  The maximum allowed number of destinations is already reached.
- **DEST_NAME_ALREADY_EXISTS**

  The destination name already exists.
- **DEST_NAME_IS_DESTGROUP_NAME**

  The destination name already exists as a destination group name.
- **EMAIL_NOT_ALLOWED_FOR_DEST_TYPE**

This type of destination cannot have an email address.

- **GATEWAY_NAME_APPEARS_TWICE**

  The gateway name appears twice on the list.
- **GATEWAY_NAME_DOES_NOT_EXIST**

  The gateway name does not exist.
- **SMSGWS_NOT_ALLOWED_FOR_DEST_TYPE**

  This type of destination cannot have SMS gateways.
- **SMTPGWS_NOT_ALLOWED_FOR_DEST_TYPE**

  This type of destination cannot have SMTP gateways.
- **SNMP_MANAGER_MUST_BE_SPECIFIED_FOR_DEST_TYPE**

  This type of destination must have an SNMP manager.
- **SNMP_MANAGER_NOT_ALLOWED_FOR_DEST_TYPE**

  This type of destination cannot have an SNMP manager.
- **NO_SMS_GATEWAYS_ARE_DEFINED**

  An SMS destination cannot be defined if no SMS gateways are defined.
- **HTTPS_ADDRESS_NOT_ALLOWED_FOR_DEST_TYPE**

  This type of destination cannot have an HTTPS address.
- **PROXY_ADDRESS_NOT_ALLOWED_FOR_DEST_TYPE**

  This type of destination cannot have a proxy address.
- **SNMP_DESTS_CANNOT_REFER_TO_USERS**

  An SNMP destination cannot refer to a user.
- **HTTPS_DESTS_CANNOT_REFER_TO_USERS**

  An HTTPS destination cannot refer to a user.
- **NO_SMTP_GATEWAYS_ARE_DEFINED**

  An email destination cannot be defined if no SMTP gateways are defined.
- **USER_EMAIL_ADDRESS_IS_NOT_DEFINED**

  The user's email address is not defined.
- **USER_PHONE_NUMBER_IS_NOT_DEFINED**

  The user's phone number is not defined.
- **USER_NAME_DOES_NOT_EXIST**

  The user name does not exist.
- **INTERNAL_DESTS_CANNOT_REFER_TO_USERS**

  An internal destination cannot refer to a user.
- **DAY_APPEARS_TWICE**

  The day '*Day*' appears twice in the list.

  **Troubleshooting:** Make sure that each day appears in the list only once.
- **HTTPS_ADDRESS_MUST_BE_SPECIFIED_FOR_DEST_TYPE**

  This type of destination must have an HTTPS address.
- **DEST_TYPE_NOT_SUPPORTED**

  This destination type is not supported.

  **Troubleshooting:** Contact IBM Support

- **USER_IS_NOT_IN_DESTINATION_DOMAINS**

  The user must be included in destination domains.
- **DOMAIN_DOESNT_EXIST**

  The domain does not exist.

## Deleting a destination

Use the **dest_delete** command to delete an event notification destination.

```
dest_delete dest=DestName
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **dest** | Object name | Name of the destination to be deleted. | Y |

Destinations that are part of a destination group or used by a rule cannot be deleted.

Destinations cannot be deleted while there are uncleared alerting events.

**Example:**

```
dest_delete dest=itmanager
```

**Output:**

```
Command completed successfully.
```

### Access control

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

### Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_DELETE_DESTINATION**

  Are you sure you want to delete destination *Destination*?

### Return codes

- **DEST_NAME_DOES_NOT_EXIST**

  The destination name does not exist.
- **CANNOT_CHANGE_EVENT_CONF_WITH_ALERTING_EVENTS**

  Cannot change an event configuration while there are alerting events.

**Troubleshooting:** Clear all alerting events before changing an event configuration.

- **DEST_IS_PART_OF_DESTGROUP**

  The destination cannot be deleted because it is part of a destination group.

- **DEST_APPEARS_IN_RULE**

  The destination cannot be deleted because it appears in a rule.

  **Troubleshooting:** To delete the destination, first delete the rule.

## Listing event notification destinations

Use the **dest_list** command to list event notification destinations.

```
dest_list [ dest=DestName ] [ type=<SNMP|EMAIL|SMS|HTTPS> ] [ domain=DomainName ] [ internal=<yes|
no> ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **dest** | Object name | Destinations to be listed. | N | All destinations. |
| **type** | Enumeration | Filter only destinations of the specified type. | N | All types. |
| **internal** | Enumeration | Filter destinations by their internal XIV attribute. | N | no |
| **domain** | Object name | The domain name. | N | All Domains |

This command lists the configuration of all defined destinations, or of a specific destination.

| Field ID | Field output | Default position |
|---|---|---|
| **name** | Name | 1 |
| **type** | Type | 2 |
| **email_address** | Email Address | 3 |
| **area_code** | Area Code | 4 |
| **number** | Phone Number | 5 |
| **snmp_manager** | SNMP Manager | 6 |
| **uri** | HTTPS Address | 7 |
| **gateways** | Gateways | N/A |
| **user** | User | 8 |
| **heartbeat_test_days** | Heartbeat Days | N/A |
| **heartbeat_test_hour** | Heartbeat Time | N/A |
| **creator** | Creator | N/A |
| **proxy** | proxy server address | N/A |
| **proxy_port** | proxy port number | N/A |

**Example:**

```
dest_list
```

**Output:**

```
Name              Type  Email Address                Phone Number  Gateways
storagemanager    EMAIL storageadmin@example.com                   all
monitoringserver  SNMP
```

### Access control

| User Category | Permission | Condition |
|---|---|---|
| Storage administrator | Conditionally Allowed | Allowed, unless the internal parameter is specified. |
| Storage integration administrator | Disallowed | N/A |
| Application administrator | Conditionally Allowed | Allowed, unless the internal parameter is specified. |
| Security administrator | Disallowed | N/A |
| Read-only users | Conditionally Allowed | Allowed, unless the internal parameter is specified. |
| Technicians | Allowed | N/A |

## Renaming a destination

Use the **dest_rename** command to rename an event notification destination.

```
dest_rename dest=DestName new_name=Name
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **dest** | Object name | The destination to be renamed. | Y |
| **new_name** | Object name | New name of the destination. | Y |

**Example:**

```
dest_rename dest=adminemail new_name=storagemanager
```

**Output:**

```
Command completed successfully.
```

### Access control

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

### Return codes

- **CANNOT_CHANGE_EVENT_CONF_WITH_ALERTING_EVENTS**

  Cannot change an event configuration while there are alerting events.

  **Troubleshooting:** Clear all alerting events before changing an event configuration.
- **DEST_NAME_DOES_NOT_EXIST**

  The destination name does not exist.
- **DEST_NAME_IS_DESTGROUP_NAME**

  The destination name already exists as a destination group name.
- **DEST_NAME_ALREADY_EXISTS**

  The destination name already exists.

## Testing a destination

Use the **dest_test** command to send a test message to an event notification destination.

```
dest_test dest=DestName management_ip=IPaddress [ smtpgw=SMTPGatewayName ]
[ smsgw=SMSGatewayName ]
 [ internal=<yes|no> ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **dest** | Object name | Name of the destination to be tested. | Y | N/A |
| **management_ip** | N/A | Management IP used for sending the event notification. | Y | N/A |
| **smtpgw** | Object name | SMTP gateway to be tested. | N | Default system choice. |
| **smsgw** | Object name | SMS gateway to be tested. | N | Default system choice. |
| **internal** | Boolean | Must be specified for XIV-internal destinations. | N | no |

This command tests a destination by sending a test message, SMS or SNMP trap. Note that a successful return code from this command does not ensure notification delivery.

Some problems with SNMP, email, and SMS delivery may fail to be detected.

For email messages, the SMTP gateway must be specified (the destination is only tested through that gateway). The same applies to the SMS gateway.

### Access control

| User Category | Permission | Condition |
|---|---|---|
| Storage administrator | Conditionally Allowed | Allowed, unless the internal parameter is specified. |
| Storage integration administrator | Disallowed | N/A |
| Application administrator | Conditionally Allowed | Allowed, unless the internal parameter is specified. |
| Security administrator | Disallowed | N/A |

| User Category | Permission | Condition |
|---|---|---|
| Read-only users | Disallowed | N/A |
| Technicians | Allowed | N/A |

### Return codes

- **DEST_NAME_DOES_NOT_EXIST**

  The destination name does not exist.
- **DEST_TEST_NOT_PERFORMED_SYSTEM_BUSY**

  The test of destination '*Destination Name*' was not performed because the system is busy.

  **Troubleshooting:** Retry in a few seconds.
- **GATEWAY_NAME_DOES_NOT_EXIST**

  The gateway name does not exist.
- **SMSGWS_MUST_BE_SPECIFIED_FOR_DEST_TYPE**

  This type of destination must have SMS gateways.
- **SMSGWS_NOT_ALLOWED_FOR_DEST_TYPE**

  This type of destination cannot have SMS gateways.
- **SMTPGWS_MUST_BE_SPECIFIED_FOR_DEST_TYPE**

  This type of destination must have SMTP gateways.
- **SMTPGWS_NOT_ALLOWED_FOR_DEST_TYPE**

  This type of destination cannot have SMTP gateways.
- **DEST_TEST_FAILED**

  The test of destination '*Destination Name*' failed.
- **SYSTEM_HAS_NO_SUCH_EXTERNAL_IP**

  The system has no such external IP address.
- **MODULE_CANNOT_SEND_MESSAGES**

  The selected module cannot send messages.

  **Troubleshooting:** Contact IBM Support
- **ONLY_TECHNICIAN_CAN_REFER_TO_INTERNAL_EVENT_OBJECTS**

  Only technicians are allowed to refer to internal event objects.

## Updating an event notification destination

Use the **dest_update** command to update a destination.

```
dest_update dest=DestName [ snmp_manager=SNMPManager ] [ uri=HTTPSaddress ] [ proxy=ProxyAddress ]
 [ proxy_port=ProxyPortNum ] [ email_address=email ]  [ smtpgws=<SMTPGW1[,SMTPGW2]...|ALL> ]
 [ area_code=AreaCode ] [ number=PhoneNumber ]  [ smsgws=<SMSGW1[,SMSGW2]...|ALL> ]
 [ user=UserName ] [ heartbeat_test_hour=HH:MM ] [ heartbeat_test_days=Day ]
 [ domain=DomainList ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **dest** | Object name | Destination name. | Y | N/A |

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **snmp_manager** | N/A | IP address or DNS name of the SNMP manager. | N | Keep unchanged. |
| **uri** | N/A | IP address or DNS name of HTTPS server. | N | Keep unchanged. |
| **proxy** | N/A | IP address or DNS name of proxy server to send HTTPS over. | N | Keep unchanged. |
| **proxy_port** | Integer | Proxy port number to send HTTPS through (1080 by default). | N | Keep unchanged. |
| **domain** | N/A | Attach the destination to the specified domains. To define more than one domain, separated them with a comma. To specify all existing domains, use "*". | N | Keep unchanged |
| **email_address** | N/A | Email address. | N | Keep unchanged. |
| **smtpgws** | Object name | List of SMTP gateways to be used. | N | Keep unchanged. |
| **area_code** | N/A | Area code of the cellular number for SMS notification. | N | Keep unchanged. |
| **number** | N/A | Cellular number for SMS notification. | N | Keep unchanged. |
| **smsgws** | Object name | SMS gateways to be used. | N | Keep unchanged. |
| **user** | Object name | User name, where the user's email or phone are used. | N | Keep unchanged. |
| **heartbeat_test_hour** | N/A | The hour of periodic heartbeat testing | N | Keep unchanged. |
| **heartbeat_test_days** | N/A | List of days for heartbeat testing: a comma-separated list of 3-letter day names (such as "mon", "mon,fri", and so on). | N | Keep unchanged. |

The parameters of this command are identical to the Defining a new event notification destination command, except that the destination type cannot be changed. All relevant fields must be specified (not only the ones that are being changed).

**Example:**

```
dest_update dest=storagemanager email_address=admin@example.com
```

**Output:**

```
Command completed successfully.
```

## Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |

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

## Return codes

- **DEST_NAME_DOES_NOT_EXIST**

  The destination name does not exist.
- **AREA_CODE_MUST_BE_SPECIFIED_FOR_DEST_TYPE**

  This type of destination must have an area code.
- **AREA_CODE_NOT_ALLOWED_FOR_DEST_TYPE**

  This type of destination cannot have an area code.
- **CANNOT_CHANGE_EVENT_CONF_WITH_ALERTING_EVENTS**

  Cannot change an event configuration while there are alerting events.

  **Troubleshooting:** Clear all alerting events before changing an event configuration.
- **EMAIL_MUST_BE_SPECIFIED_FOR_DEST_TYPE**

  This type of destination must have an email address.
- **EMAIL_NOT_ALLOWED_FOR_DEST_TYPE**

  This type of destination cannot have an email address.
- **GATEWAY_NAME_APPEARS_TWICE**

  The gateway name appears twice on the list.
- **GATEWAY_NAME_DOES_NOT_EXIST**

  The gateway name does not exist.
- **NUMBER_MUST_BE_SPECIFIED_FOR_DEST_TYPE**

  This type of destination must have a number.
- **NUMBER_NOT_ALLOWED_FOR_DEST_TYPE**

  This type of destination cannot have a number.
- **SMSGWS_NOT_ALLOWED_FOR_DEST_TYPE**

  This type of destination cannot have SMS gateways.
- **SNMP_MANAGER_NOT_ALLOWED_FOR_DEST_TYPE**

  This type of destination cannot have an SNMP manager.
- **NO_SMTP_GATEWAYS_ARE_DEFINED**

  An email destination cannot be defined if no SMTP gateways are defined.
- **DEST_CANNOT_HAVE_A_USER_AND_AN_EMAIL_ADDRESS**

  The destination cannot simultaneously have an email address and refer to a user.
- **DEST_CANNOT_HAVE_A_USER_AND_A_PHONE_NUMBER**

  The destination cannot simultaneously have a phone number and refer to a user.
- **USER_PHONE_NUMBER_IS_NOT_DEFINED**

  The user's phone number is not defined.

- **USER_NAME_DOES_NOT_EXIST**

  The user name does not exist.
- **INTERNAL_DESTS_CANNOT_REFER_TO_USERS**

  An internal destination cannot refer to a user.
- **DEST_HEARTBEAT_DAYS_BUT_NO_HOUR**

  Destination heartbeat days are specified, but not heartbeat hour.
- **HTTPS_ADDRESS_NOT_ALLOWED_FOR_DEST_TYPE**

  This type of destination cannot have an HTTPS address.
- **PROXY_ADDRESS_NOT_ALLOWED_FOR_DEST_TYPE**

  This type of destination cannot have a proxy address.
- **SNMP_DESTS_CANNOT_REFER_TO_USERS**

  An SNMP destination cannot refer to a user.
- **HTTPS_DESTS_CANNOT_REFER_TO_USERS**

  An HTTPS destination cannot refer to a user.
- **USER_EMAIL_ADDRESS_IS_NOT_DEFINED**

  The user's email address is not defined.
- **SMTPGWS_NOT_ALLOWED_FOR_DEST_TYPE**

  This type of destination cannot have SMTP gateways.
- **DAY_APPEARS_TWICE**

  The day '*Day*' appears twice in the list.

  **Troubleshooting:** Make sure that each day appears in the list only once.
- **SNMP_MANAGER_MUST_BE_SPECIFIED_FOR_DEST_TYPE**

  This type of destination must have an SNMP manager.
- **NO_SMS_GATEWAYS_ARE_DEFINED**

  An SMS destination cannot be defined if no SMS gateways are defined.
- **HTTPS_ADDRESS_MUST_BE_SPECIFIED_FOR_DEST_TYPE**

  This type of destination must have an HTTPS address.
- **DEST_TYPE_NOT_SUPPORTED**

  This destination type is not supported.

  **Troubleshooting:** Contact IBM Support
- **DOMAIN_DOESNT_EXIST**

  The domain does not exist.
- **USER_IS_NOT_IN_DESTINATION_DOMAINS**

  The user must be included in destination domains.
- **DESTINATION_IS_NOT_IN_RULE_DOMAINS**

  The destination must be included in rule domains.
- **DESTINATION_IS_NOT_IN_DESTGROUP_DOMAINS**

  The destination must be included in the destination group domains.

# Adding a destination to a destination group

Use the **destgroup_add_dest** command to add an event notification destination to a destination group.

```
destgroup_add_dest destgroup=GroupName dest=DestName
```

## Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **destgroup** | Object name | Destination group name to which to add the destination. | Y |
| **dest** | Object name | Destination to be added to the group. | Y |

The command fails if the destination group already contains the destination.

The command cannot be executed while there are uncleared alerting events.

**Example:**

```
destgroup_add_dest destgroup=alladmins dest=john
```

**Output:**

```
Command completed successfully.
```

## Access control

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

## Return codes

- **DESTGROUP_NAME_DOES_NOT_EXIST**

  The destination group name does not exist.
- **CANNOT_CHANGE_EVENT_CONF_WITH_ALERTING_EVENTS**

  Cannot change an event configuration while there are alerting events.

  **Troubleshooting:** Clear all alerting events before changing an event configuration.
- **DEST_NAME_DOES_NOT_EXIST**

  The destination name does not exist.
- **DESTGROUP_MAX_DESTS_REACHED**

  The maximum allowed number of destinations is already reached in destination groups.
- **DESTGROUP_ALREADY_INCLUDES_DEST**

  The destination group already includes this destination name.

- **DESTINATION_IS_NOT_IN_DESTGROUP_DOMAINS**

  The destination must be included in the destination group domains.

## Creating a destination group

Use the **destgroup_create** command to create an event notification destinations group.

```
destgroup_create destgroup=GroupName [ domain=DomainList ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **destgroup** | Object name | Destination group name. | Y | N/A |
| **domain** | N/A | Attach the destination group to the specified domains. To define more than one domain, separated them with a comma. To specify all existing domains, use "*". | N | none |

This command creates a destination group, which is used by rules to send notifications to the entire group without specifying all the destinations for each rule. You can also add or remove destinations from the group, which eliminates the need to change the configuration of each rule separately.

Upon creation, the destination group is empty. To add a destination to a destination group, use the Adding a destination to a destination group command.

**Example:**

```
destgroup_create destgroup=alladmins
```

**Output:**

```
Command completed successfully.
```

### Access control

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

### Return codes

- **DESTGROUP_MAX_REACHED**

  The maximum allowed number of destination groups is already reached.
- **DESTGROUP_NAME_ALREADY_EXISTS**

The destination group name already exists.

- **DESTGROUP_NAME_IS_DEST_NAME**

  The destination group name already exists as a destination name.
- **DOMAIN_DOESNT_EXIST**

  The domain does not exist.

## Updating an event notification destination group

Use the **destgroup_update** command to update a destination group.

```
destgroup_update destgroup=GroupName domain=DomainList
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **destgroup** | Object name | Destination group name. | Y |
| **domain** | N/A | Attach the destination group to the specified domains. To define more than one domain, separated them with a comma. To specify all existing domains, use "*". | Y |

**Example:**

```
destgroup_update destgroup=alladmins domain=D1,D2
```

**Output:**

```
Command completed successfully.
```

### Access control

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

### Return codes

- **DESTINATION_IS_NOT_IN_DESTGROUP_DOMAINS**

  The destination must be included in the destination group domains.
- **DOMAIN_DOESNT_EXIST**

  The domain does not exist.
- **DESTGROUP_IS_NOT_IN_RULE_DOMAINS**

  The destination groups must be included in rule domains.
- **DESTGROUP_NAME_DOES_NOT_EXIST**

The destination group name does not exist.

- **INTERNAL_EVENT_OBJECTS_CANNOT_USE_SPECIFIC_DOMAINS**

  Internal event objects cannot be defined on specific domains.

## Deleting a destination group

Use the **destgroup_delete** command to delete an event notification destination group.

```
destgroup_delete destgroup=GroupName
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **destgroup** | Object name | The name of the destination group to be deleted. | Y |

**Example:**

```
destgroup_delete destgroup=dest1
```

**Output:**

```
Command completed successfully
```

### Access control

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

### Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_DELETE_DESTINATION_GROUP**

  Are you sure you want to delete destination group *Destination Group*?

### Return codes

- **CANNOT_CHANGE_EVENT_CONF_WITH_ALERTING_EVENTS**

  Cannot change an event configuration while there are alerting events.

  **Troubleshooting:** Clear all alerting events before changing an event configuration.
- **DESTGROUP_APPEARS_IN_RULE**

  The destination group appears in a rule.

  **Troubleshooting:** To delete the destination group, first delete the rule.
- **DESTGROUP_NAME_DOES_NOT_EXIST**

  The destination group name does not exist.

# Listing destination groups

Use the **destgroup_list** command to list destination groups.

```
destgroup_list [ destgroup=GroupName ] [ domain=DomainName ]
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **destgroup** | Object name | Destination group to be listed. | N | All groups. |
| **domain** | Object name | The domain name. | N | All Domains |

This command lists all destination groups or a specific one. All the destinations are listed for each destination group.

| Field ID | Field output | Default position |
|---|---|---|
| **name** | Name | 1 |
| **dests** | Destinations | 2 |
| **creator** | Creator | N/A |

**Example:**

```
destgroup_list
```

**Output:**

```
Name     Destinations
itstaff  john,michael,linda,monitoringserver
```

## Access control

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

# Removing a destination from a destination group

Use the **destgroup_remove_dest** command to remove an event notification destination from a destination group.

```
destgroup_remove_dest destgroup=GroupName dest=DestName
```

## Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **destgroup** | Object name | Group name. | Y |

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **dest** | Object name | Destination to be removed from the group. | Y |

This command cannot be executed while there are uncleared alerting events.

**Example:**

```
destgroup_remove_dest destgroup=alladmins dest=john
```

**Output:**

```
Command completed successfully.
```

### Access control

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

### Return codes

- **DESTGROUP_NAME_DOES_NOT_EXIST**

  The destination group name does not exist.
- **DEST_NAME_DOES_NOT_EXIST**

  The destination name does not exist.
- **CANNOT_CHANGE_EVENT_CONF_WITH_ALERTING_EVENTS**

  Cannot change an event configuration while there are alerting events.

  **Troubleshooting:** Clear all alerting events before changing an event configuration.
- **DESTGROUP_DOES_NOT_INCLUDE_DEST**

  The destination group does not include this destination name.

## Renaming a destination group

Use the **destgroup_rename** command to rename an event notification destination group.

```
destgroup_rename destgroup=GroupName new_name=Name
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **destgroup** | Object name | Destination group to be renamed. | Y |
| **new_name** | Object name | New name of the destination group. | Y |

This command cannot be executed while there are uncleared alerting events.

**Example:**

```
destgroup_rename destgroup=alladmins new_name=itstaff
```

**Output:**

```
Command completed successfully.
```

### Access control

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

### Return codes

- **DESTGROUP_NAME_DOES_NOT_EXIST**

  The destination group name does not exist.
- **DESTGROUP_NAME_ALREADY_EXISTS**

  The destination group name already exists.
- **CANNOT_CHANGE_EVENT_CONF_WITH_ALERTING_EVENTS**

  Cannot change an event configuration while there are alerting events.

  **Troubleshooting:** Clear all alerting events before changing an event configuration.
- **DESTGROUP_NAME_IS_DEST_NAME**

  The destination group name already exists as a destination name.

## Clearing alerting events

Use the **event_clear** command to clear alerting events.

```
event_clear event_id=EventId [ all_preceding=<yes|no> ] [ internal=<yes|no|all> ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **event_id** | Positive integer | The ID number of the event to be cleared. | Y | N/A |
| **all_preceding** | Boolean | Clears all events preceding the specified event. | N | no |
| **internal** | Boolean | Clears XIV-internal events. | N | no |

In order to ensure that an event was indeed received, an event notification may be sent repeatedly until it is cleared with a CLI command or from the GUI. Such events are called *alerting* events. An event is defined as *alerting* if at the time of the event's generation it was matched by an *alerting* rule, meaning a rule that has either snooze or escalation definitions.

Notifications for the alerting event are sent until it is cleared by this command. The clearing operation does not imply that the problem has been solved. It only implies that the event has been noted by the relevant person who takes responsibility for fixing the problem.

The user may clear either a specific event or all alerting events.

**Example:**

```
event_clear event_id=87
```

**Output:**

```
Command completed successfully.
```

### Access control

| User Category | Permission | Condition |
|---|---|---|
| Storage administrator | Conditionally Allowed | Allowed, unless the internal parameter is specified. |
| Storage integration administrator | Disallowed | N/A |
| Application administrator | Conditionally Allowed | Allowed, unless the internal parameter is specified. |
| Security administrator | Disallowed | N/A |
| Read-only users | Disallowed | N/A |
| Technicians | Allowed | N/A |

### Return codes

- **ONLY_TECHNICIAN_CAN_REFER_TO_INTERNAL_EVENT_OBJECTS**

  Only technicians are allowed to refer to internal event objects.

## Listing events

Use the **event_list** command to list system events.

```
event_list [ max_events=MaxEventsToList ] [ after=TimeStamp ] [ before=TimeStamp ]
 [ min_severity=<INFORMATIONAL|WARNING|MINOR|MAJOR|CRITICAL> ] [ alerting=<yes|no|all> ]
 [ cleared=<yes|no|all> ] [ code=EventCode ]
 [ object_type=<cons_group|destgroup|dest|dm|host|map|mirror|pool|rule|smsgw|smtpgw|target|v
olume|cluster|ip_interface|ldap_conf|meta_data_object|sync_schedule|user|user_group|ldap_ser
ver|modules_status|elicense|ipsec_connection|ipsec_tunnel|cross_cons_group,...> ]
 [ internal=<yes|no|all> ] [ beg=BeginIndex ] [ end=EndIndex ] [ count_all=<yes|no> ]
 [ domain=DomainName ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **max_events** | Positive integer | Maximal number of events to list (see below for more details). | N | 300 |
| **after** | N/A | Earliest time/date. | N | No filter. |

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **before** | N/A | Latest time/date. | N | No filter. |
| **min_severity** | Enumeration | Minimum severity. | N | No filter. |
| **alerting** | Boolean | Filter alerting events. | N | No filter. |
| **cleared** | Boolean | Filter cleared events. | N | No filter. |
| **code** | N/A | Filter by a specific event code. | N | No filter. |
| **object_type** | Enumeration | Filter events by the type of the related system object. | N | No filter. |
| **internal** | Boolean | Filter XIV internal events. | N | No filter. |
| **beg** | Integer | Index of the first event to list. If negative, then counts from the end. | N | 1 |
| **end** | Integer | Index of the last event to list (not inclusive). If negative, then counts from the end. | N | Last event + 1. |
| **count_all** | Boolean | If yes, it scans all the events between beginning and end for computing the number of events meeting the criteria. | N | No. |
| **domain** | Object name | The domain name. | N | All Domains. |

This command lists system events according to specified criteria, such as minimum severity, event type, and so on. The event list displays the following information for each event: timestamp, severity, code, user and description.

The max_events parameter

Events are listed and sorted by the time of creation in ascending order.

If the maximal number of events to display is limited by specifying the **max_events** parameter, and the number of retrieved events exceeds the maximum, only the most recent events will be displayed.

If the **max_events** parameter is not specified, the number of displayed events defaults to 300.

The fields `before` and `after`

The syntax for the `before` and `after` fields is as follows: `Y-M-D[.[h[:m[:s]]]]`, where the ranges are as follows:

- Y - year (four digit)
- M - month (1-12)
- D - day (1-31)
- h - hour (0-23, with 0 as default)
- m - minute (0-59, with 0 as default)
- s - second (0-59, with 0 as default)

The year, month and day are separated by dashes, while the optional hour, minute and second are separated by colons.

**Example:**

```
event_list max_events=10
```

**Output:**

```
Timestamp             Severity        Code
2009-05-12 15:10:16   Informational   START_WORK
2009-05-12 15:16:11   Informational   POOL_CREATE
2009-05-12 15:16:23   Informational   VOLUME_CREATE

Additional output fields
(lines are broken to fit the page width of this Guide):

User              Description
                  System has entered ON state.
xiv_development   Storage Pool of size 171GB was created with name
                  'p1_m'.
xiv_development   Volume was created with name 'master' and size 17GB in
                  Storage Pool with name 'p1_m'.
```

| Field ID | Field output | Default position |
|---|---|---|
| **timestamp** | Timestamp | 1 |
| **severity** | Severity | 2 |
| **code** | Code | 3 |
| **user_name** | User | 4 |
| **description** | Description | 5 |
| **index** | Index | N/A |
| **alerting** | Alerting | N/A |
| **cleared** | Cleared | N/A |
| **tshooting** | Trouble Shooting | N/A |

## Access control

| User Category | Permission | Condition |
|---|---|---|
| Storage administrator | Conditionally Allowed | Allowed, unless the internal parameter is specified. |
| Storage integration administrator | Conditionally Allowed | Allowed, unless the internal parameter is specified. |
| Application administrator | Conditionally Allowed | Allowed, unless the internal parameter is specified. |
| Security administrator | Conditionally Allowed | Allowed, unless the internal parameter is specified. |
| Read-only users | Conditionally Allowed | Allowed, unless the internal parameter is specified. |
| Technicians | Allowed | N/A |

## Return codes

- **UNRECOGNIZED_EVENT_CODE**

  '*String*' is not a recognized return code.

  **Troubleshooting:** Consult the manual for the list of valid return codes.
- **CANNOT_READ_EVENTS**

  Cannot read events.

  **Troubleshooting:** Contact IBM Support.
- **DOMAIN_DOESNT_EXIST**

The domain does not exist.

## Listing uncleared alerting events

Use the **event_list_uncleared** command to list uncleared alerting events.

```
event_list_uncleared [ domain=DomainName ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **domain** | Object name | The domain name. | N | All Domains |

**Example:**

```
event_list_uncleared
```

**Output:**

```
Index   Code            Severity
------- --------------- ---------------
318     VOLUME_CREATE   Informational
666     VOLUME_DELETE   Informational
```

| Field ID | Field output | Default position |
|---|---|---|
| **index** | Index | 1 |
| **code** | Code | 2 |
| **severity** | Severity | 3 |

### Access control

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

## Setting the threshold for event notification

Use the **event_redefine_threshold** command to redefine the threshold of a parameterized event.

```
event_redefine_threshold code=EventCode severity=<INFORMATIONAL|WARNING|MINOR|MAJOR|CRITICAL|
NONE> threshold=<
ThresholdValue|NONE>
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **code** | N/A | Event code. | Y |
| **severity** | Enumeration | Severity. | Y |
| **threshold** | Integer | Threshold value, or NONE to indicate that an event with this severity is not created. | Y |

This command can be applied to parameterized events, that is events that are triggered when a certain parameter crosses a certain threshold. Using this command the user can change the threshold for event notification. Furthermore, multiple thresholds can be defined using multiple invocations of this command, one for each event severity. When the relevant parameter crosses a threshold, an event with the matching severity is created.

### Access control

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

### Return codes

- **EVENT_DOES_NOT_HAVE_THRESHOLDS**

  The event does not have thresholds.
- **EVENT_THRESHOLD_IS_ILLEGAL**

  An illegal value for the event threshold.

  **Troubleshooting:** Event threshold values must be monotonic.
- **UNRECOGNIZED_EVENT_CODE**

  '*String*' is not a recognized return code.

  **Troubleshooting:** Consult the manual for the list of valid return codes.
- **LAST_EVENT_THRESHOLD_CANNOT_BE_DELETED**

  The event must have at least one threshold value.

## Listing thresholds

Use the **event_threshold_list** to list event thresholds.

```
event_threshold_list [ code=EventCode ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **code** | Enumeration | Filter by a specific event code. | N | no filter |

| Field ID | Field output | Default position |
|---|---|---|
| **code** | Code | 1 |
| **has_thresholds** | Has Thresholds? | N/A |
| **not_in_use** | Not In Use | N/A |
| **replaced_by** | Replaced By | N/A |
| **default_thresholds.0** | INFORMATIONAL(def) | 7 |
| **default_thresholds.1** | WARNING(def) | 8 |
| **default_thresholds.2** | MINOR(def) | 9 |
| **default_thresholds.3** | MAJOR(def) | 10 |
| **default_thresholds.4** | CRITICAL(def) | 11 |
| **thresholds.0** | INFORMATIONAL | 2 |
| **thresholds.1** | WARNING | 3 |
| **thresholds.2** | MINOR | 4 |
| **thresholds.3** | MAJOR | 5 |
| **thresholds.4** | CRITICAL | 6 |

**Example:**

```
event_threshold_list
```

**Output:**

```
Code                                    INFORMATIONAL   WARNING   MINOR
--------------------------------------- --------------- --------- -------
STORAGE_POOL_SNAPSHOT_USAGE_INCREASED   none             80        90
STORAGE_POOL_VOLUME_USAGE_INCREASED     none             80        90

MAJOR   CRITICAL   INFORMATIONAL(def)   WARNING(def)   MINOR(def)   MAJOR(def)
------- ---------- -------------------- -------------- ------------ ------------
95      none       none                 80             90            95
95      none       none                 80             90            95

CRITICAL(def)
---------------
none
none
```

## Access control

| User Category | Permission | Condition |
|---|---|---|
| Storage administrator | Conditionally Allowed | Allowed, unless the internal parameter is specified. |
| Storage integration administrator | Disallowed | N/A |
| Application administrator | Conditionally Allowed | Allowed, unless the internal parameter is specified. |
| Security administrator | Disallowed | N/A |
| Read-only users | Conditionally Allowed | Allowed, unless the internal parameter is specified. |
| Technicians | Allowed | N/A |

# Activating a rule

Use the **rule_activate** command to activate an event notification rule.

```
rule_activate rule=RuleName
```

## Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **rule** | Object name | The name of the rule to be activated. | Y |

This command activates the specified rule. An active rule is matched against events and generates notifications. If the rule is already active, this command has no effect.

## Access control

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

## Return codes

- **EVENT_RULE_NAME_DOES_NOT_EXIST**

  The event rule name does not exist.

# Creating event notification rules

Use the **rule_create** command to create an event notification rule.

```
rule_create rule=RuleName [ min_severity=<INFORMATIONAL|WARNING|MINOR|MAJOR|CRITICAL|NONE> ]
 [ codes=Codes | except_codes=EventCodes ] [ escalation_only=<yes|no> ]
 dests=dest1,dest2,... [ snooze_time=SnoozeTime ]
 [ escalation_time=EscalationTime escalation_rule=EscalationRule ] [ domain=DomainList ]
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **rule** | Object name | The name of the new rule. | Y | N/A |
| **min_severity** | Enumeration | Minimal event severity for rule filtering. | N | All severities. |
| **codes** | N/A | Filter only events with these codes. | N | All events. |
| **except_codes** | N/A | Filter only events with other codes. | N | All events. |
| **escalation_only** | Boolean | Specifies that this rule can only be used for escalation. | N | no |

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **dests** | Object name | Comma-separated list of destinations and destination groups for event notification. | Y | N/A |
| **snooze_time** | Integer | Snooze time in minutes. | N | No snoozing. |
| **escalation_rule** | Object name | Escalation rule. | N | N/A |
| **escalation_time** | Integer | Escalation time in minutes. Escalation time should not be smaller than snooze time. Refer to escalation_rule above for more information. | N | No escalation. |
| **domain** | N/A | The rule will be attached to the specified domains. To define more than one domain, separate them with a comma. To specify all existing domains, use "*". | N | none |

This command defines a new event notification rule. An event notification rule determines which events should generate which notifications. When an event occurs, it is checked by all currently defined rules, based on which notifications are generated.

Each rule has a filtering and notification configuration.

The filtering configuration controls which events match this rule. The filtering can be based on the event's code, by specifying a minimum severity. When using this configuration, each event with a severity higher or equal to the rule's **min_severity** parameter matches this rule. Alternatively, the rule may match only a specific event code. Two filters can be combined for events whose severity depends on a run-time parameter.

The second part of a rule configuration is a list of destinations and destination groups that receive the notification when an event matches the filtering criteria. If a destination is included both in the rule and in one of the rule's destination groups, it still gets only one notification. The same applies if a destination is included in two destination groups, or if the event matches the filtering criteria of several rules, all using the same destination.

A rule can be defined as *alerting*, which means that notifications are sent repeatedly until the matching events are cleared using the Clearing alerting events command.

Clearing the event does not mean that the problem has been solved. It only means that it was noticed and there is no need to continue sending notifications.

The repeated sending of notifications can be defined by two ways:

- The **snooze** parameter causes the notifications to be sent again and again to the same destinations. The time in minutes between the repeated transmissions is determined by the `snooze` value.
- The **escalation_time** and **escalation_rule** parameters cause the notifications to be sent to the destination list of the **escalation_rule** if it is not cleared within `escalation_time` minutes.

Rules can escalate only to alerting rules (that is, to rules that have snooze or escalation definitions) in order to prevent a situation where notifications are stopped from being sent.

A rule cannot escalate to itself, nor can it be defined in a cyclic escalation of rules.

The **escalation_only** parameter defines a rule without filters, which can only be used as an escalation for other rules.

The snooze time cannot be greater than the escalation time.

It is not permitted to define new rules while there are uncleared alerting events.

The following example sends alerts upon critical events to John's cellular number and to the emails of all the IT staff. The alerts will be resent every 20 minutes until the events are cleared.

**Example:**

```
rule_create rule=critical_alerts min_severity=critical destinations=john-cell,itstaff
snooze_time=20
```

**Output:**

```
Command completed successfully.
```

## Access control

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

## Return codes

- **EVENT_RULE_MAX_REACHED**

  The maximum allowed number of event rules is already reached.
- **EVENT_RULE_CANNOT_ESCALATE_TO_NON_ALERTING_RULES**

  An event rule cannot be escalated to a non-alerting rule.

  **Troubleshooting:** An alerting rule can only be escalated to another escalation rule.
- **DEST_APPEARS_TWICE**

  The destination or destination group appears twice.
- **EVENT_RULE_NAME_ALREADY_EXISTS**

  The event rule name already exists.
- **EVENT_RULE_NAME_DOES_NOT_EXIST**

  The event rule name does not exist.
- **NAME_IS_NEITHER_DEST_NOR_GROUP**

  The specified ame is neither a destination group name nor a destination name.
- **ESCALATION_TIME_MUST_BE_LARGER_THAN_SNOOZE_TIME**

  Escalation time must be larger than snooze time.
- **RULE_MAX_DESTS_REACHED**

  The maximum allowed number of destinations and destination groups in a rule is already reached.
- **EVENT_RULE_MUST_HAVE_FILTER**

  An alerting event rule must have a filter represented by an event code or severity.
- **EVENT_RULE_CANNOT_REFER_TO_INTERNAL_EVENT_CODES**

  A user event rule cannot refer to internal event codes.

- **ESCALATION_EVENT_RULE_CANNOT_HAVE_FILTER**

  An escalation-only event rule cannot have code or min_severity specification.
- **ESCALATION_EVENT_RULE_MUST_BE_ALERTING**

  An escalation-only event rule must be an alerting rule.
- **TOO_MANY_EVENT_CODES**

  A maximum of *Maximum* return codes can be defined.
- **EVENT_CODE_APPEARS_TWICE**

  The return code '*Code*' appears twice in the list.

  **Troubleshooting:** Make sure that each return code appears in the list only once.
- **UNRECOGNIZED_EVENT_CODE**

  '*String*' is not a recognized return code.

  **Troubleshooting:** Consult the manual for the list of valid return codes.
- **EVENT_RULE_CANNOT_HAVE_A_CATEGORY**

  A user event rule cannot have a category definition.
- **DOMAIN_DOESNT_EXIST**

  The domain does not exist.
- **DESTINATION_IS_NOT_IN_RULE_DOMAINS**

  The destination must be included in rule domains.
- **DESTGROUP_IS_NOT_IN_RULE_DOMAINS**

  The destination groups must be included in rule domains.
- **ESCALATION_RULE_NOT_IN_RULE_DOMAINS**

  An escalation rule must belong to rule domains.
- **EVENT_RULE_MUST_NOT_HAVE_SNMP_DEST**

  According to the current system configuration state (snmp_type = NONE), an event rule must not have an SNMP destination.

# Deactivating a rule

Use the **rule_deactivate** command to deactivate an event notification rule.

```
rule_deactivate rule=RuleName
```

## Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **rule** | Object name | The name of the rule to be deactivated. | Y |

A deactivated rule is not matched against events and does not generate notifications. If the rule is already inactive, then this command has no effect.

Inactive rules cannot be used as escalation rules.

The rules of type `escalation_only` cannot be deactivated.

### Access control

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

### Return codes

- **EVENT_RULE_NAME_DOES_NOT_EXIST**

  The event rule name does not exist.
- **ESCALATION_ONLY_RULES_ARE_ALWAYS_ACTIVE**

  An escalation-only event rule cannot be deactivated or activated.

## Deleting event notification rules

Use the **rule_delete** command to delete an event notification rule.

```
rule_delete rule=RuleName
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **rule** | Object name | The rule to be deleted. | Y |

Rules that are defined as the escalation of other rules cannot be deleted.

It is not permitted to delete a rule while there are uncleared alerting events.

**Example:**

```
rule_delete rule=emergency_alerts
```

**Output:**

```
Command completed successfully.
```

### Access control

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

### Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_DELETE_RULE**

  Are you sure you want to delete rule *Rule*?

### Return codes

- **CANNOT_CHANGE_EVENT_CONF_WITH_ALERTING_EVENTS**

  Cannot change an event configuration while there are alerting events.

  **Troubleshooting:** Clear all alerting events before changing an event configuration.
- **EVENT_RULE_NAME_DOES_NOT_EXIST**

  The event rule name does not exist.
- **EVENT_RULE_USED_FOR_ESCALATION_CAN_NOT_BE_DELETED**

  The event rule is an escalation rule of another event rule. Therefore, it cannot be deleted.

  **Troubleshooting:** Delete all escalation rules that refer to this rule as their escalation rule.

## Listing event notification rules

Use the **rule_list** command to list event notification rules.

```
rule_list [ rule=RuleName ] [ internal=<yes|no> ] [ domain=DomainName ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **rule** | Object name | The rule to be listed. | N | All rules. |
| **internal** | Enumeration | Filters XIV internal rules. | N | no |
| **domain** | Object name | The domain name. | N | All Domains |

**Example:**

```
rule_list
```

**Output:**

```
Name             Minimum Severity  Event Code  Destinations
---------------  ----------------  ----------  -----------------
emergency_alerts  critical          all         john-cell,itstaff
```

| Field ID | Field output | Default position |
|---|---|---|
| **name** | Name | 1 |
| **min_severity** | Minimum Severity | 2 |
| **codes** | Event Codes | 3 |
| **except_codes** | Except Codes | 4 |
| **dests** | Destinations | 5 |
| **active** | Active | 6 |
| **escalation_time** | Escalation Time | N/A |
| **snooze_time** | Snooze Time | N/A |

| Field ID | Field output | Default position |
|---|---|---|
| **escalation_rule** | Escalation Rule | N/A |
| **escalation_only** | Escalation Only | 7 |
| **category** | Category | N/A |
| **creator** | Creator | N/A |

### Access control

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

## Renaming event notification rules

Use the **rule_rename** command to rename an event notification rule.

```
rule_rename rule=RuleName new_name=Name
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **rule** | Object name | The rule to be renamed. | Y |
| **new_name** | Object name | The new name of the rule. | Y |

**Example:**

```
rule_rename rule=critical_alerts new_name=emergency_alerts
```

**Output:**

```
Command completed successfully
```

### Access control

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

### Return codes

- **CANNOT_CHANGE_EVENT_CONF_WITH_ALERTING_EVENTS**

Cannot change an event configuration while there are alerting events.

**Troubleshooting:** Clear all alerting events before changing an event configuration.

- **EVENT_RULE_NAME_ALREADY_EXISTS**

  The event rule name already exists.

- **EVENT_RULE_NAME_DOES_NOT_EXIST**

  The event rule name does not exist.

## Updating an event notification rule

Use the **rule_update** command to update an event notification rule.

```
rule_update rule=RuleName [ min_severity=<INFORMATIONAL|WARNING|MINOR|MAJOR|CRITICAL|NONE> ]
[ codes=Codes ]
 [ except_codes=EventCodes ] [ escalation_only=<yes|no> ] [ dests=dest1,dest2,... ]
 [ snooze_time=SnoozeTime ] [ escalation_time=EscalationTime ]
 [ escalation_rule=EscalationRule ] [ domain=DomainList ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **rule** | Object name | The name of the rule. | Y | N/A |
| **min_severity** | Enumeration | Minimum event severity for rule filtering. | N | Leave unchanged. |
| **codes** | N/A | Filter only events with this code. | N | Leave unchanged. |
| **except_codes** | N/A | Filter only events with other codes. | N | Leave unchanged. |
| **escalation_only** | Boolean | Specifies that this rule can only be used for escalation. | N | no |
| **dests** | Object name | Comma-separated list of destinations and destination groups for event notification. | N | Leave unchanged. |
| **snooze_time** | Integer | Snooze time in minutes. | N | Leave unchanged. |
| **escalation_time** | Integer | Escalation time in minutes. | N | Leave unchanged. |
| **escalation_rule** | Object name | Escalation rule. | N | Leave unchanged. |
| **domain** | N/A | The rule will be attached to the specified domains. To specify more than one domain, separate them with a comma. To specify all existing domains, use "*". | N | Leave unchanged. |

This command updates the configuration of an event notification rule. All parameters and their descriptions are identical to the Creating event notification rules command.

Parameters which are not specified are not changed.

**Example:**

```
rule_update rule=critical_alerts min_severity=critical destinations=john-cell,itstaff
snooze_time=30
```

**Output:**

```
Command completed successfully.
```

## Access control

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

## Return codes

- **EVENT_RULE_NAME_DOES_NOT_EXIST**

  The event rule name does not exist.
- **CANNOT_CHANGE_EVENT_CONF_WITH_ALERTING_EVENTS**

  Cannot change an event configuration while there are alerting events.

  **Troubleshooting:** Clear all alerting events before changing an event configuration.
- **EVENT_RULE_CANNOT_ESCALATE_TO_ITSELF**

  An event rule cannot be its own escalation rule.
- **EVENT_RULE_CANNOT_ESCALATE_TO_NON_ALERTING_RULES**

  An event rule cannot be escalated to a non-alerting rule.

  **Troubleshooting:** An alerting rule can only be escalated to another escalation rule.
- **DEST_APPEARS_TWICE**

  The destination or destination group appears twice.
- **EVENT_RULE_MISSING_ESCALATION_RULE**

  An alerting event rule must have an escalation rule.

  **Troubleshooting:** If escalation time is specified, then an escalation rule must be specified as well.
- **EVENT_RULE_MISSING_ESCALATION_TIME**

  An alerting event rule must have escalation time.

  **Troubleshooting:** If an escalation rule is specified, then escalation time must be specified also.
- **NAME_IS_NEITHER_DEST_NOR_GROUP**

  The specified ame is neither a destination group name nor a destination name.
- **ESCALATION_TIME_MUST_BE_LARGER_THAN_SNOOZE_TIME**

  Escalation time must be larger than snooze time.
- **RULE_MAX_DESTS_REACHED**

  The maximum allowed number of destinations and destination groups in a rule is already reached.
- **EVENT_RULE_MUST_HAVE_FILTER**

  An alerting event rule must have a filter represented by an event code or severity.
- **CYCLIC_ESCALATION_RULES_DEFINITION**

Event rule escalation cannot be cyclic.

- **EVENT_RULE_USED_FOR_ESCALATION_MUST_BE_ALERTING**

  The event rule is an escalation rule of another event rule, and thus must be an alerting rule.
- **EVENT_RULE_CANNOT_REFER_TO_INTERNAL_EVENT_CODES**

  A user event rule cannot refer to internal event codes.
- **ESCALATION_EVENT_RULE_CANNOT_HAVE_FILTER**

  An escalation-only event rule cannot have code or min_severity specification.
- **EVENT_RULE_CANNOT_HAVE_A_CATEGORY**

  A user event rule cannot have a category definition.
- **EVENT_RULE_CANNOT_HAVE_BOTH_CODES_AND_EXCEPTION_CODES**

  An event rule cannot have both codes and exception codes.
- **ESCALATION_EVENT_RULE_MUST_BE_ALERTING**

  An escalation-only event rule must be an alerting rule.
- **TOO_MANY_EVENT_CODES**

  A maximum of *Maximum* return codes can be defined.
- **EVENT_CODE_APPEARS_TWICE**

  The return code '*Code*' appears twice in the list.

  **Troubleshooting:** Make sure that each return code appears in the list only once.
- **UNRECOGNIZED_EVENT_CODE**

  '*String*' is not a recognized return code.

  **Troubleshooting:** Consult the manual for the list of valid return codes.
- **DOMAIN_DOESNT_EXIST**

  The domain does not exist.
- **DESTINATION_IS_NOT_IN_RULE_DOMAINS**

  The destination must be included in rule domains.
- **DESTGROUP_IS_NOT_IN_RULE_DOMAINS**

  The destination groups must be included in rule domains.
- **ESCALATION_RULE_NOT_IN_RULE_DOMAINS**

  An escalation rule must belong to rule domains.
- **EVENT_RULE_MUST_NOT_HAVE_SNMP_DEST**

  According to the current system configuration state (snmp_type = NONE), an event rule must not have an SNMP destination.

## Defining an SMS gateway

Use the **smsgw_define** command to define an SMS gateway.

```
smsgw_define smsgw=SMSGatewayName email_address=email subject_line=SubjectLineScheme
email_body=EmailBodySchem
e [ smtpgw=<SMTPGW1[,SMTPGW2]...|ALL> ]
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **smsgw** | Object name | SMS gateway name. | Y | N/A |
| **email_address** | Token String | Format for the email address. | Y | N/A |
| **subject_line** | Token String | Format for the subject line. | Y | N/A |
| **email_body** | Token String | Format for the email body. | Y | N/A |
| **smtpgw** | Object name | List of SMTP gateways to be used. | N | The SMTP gateways defined in the smtpgw_prioritize command. |

SMS gateways are used to send event notifications via SMS messages. SMS messages are sent via SMS-to-email servers. To definw a new SMS gateway, it is necessary to know how SMS messages are encapsulated in the email message.

When the system sends an SMS message, it uses the actual message text that describes the event and the destination number. The destination number is comprised from an area code and the local number. Both are specified when a destination is defined as described in the Defining a new event notification destination command.

The message text and the destination numbers can be embedded into the email message into various parts of the email message: destination address, subject line, or email body. This command defines how email messages are formatted, and how they contain the information of the specific SMS.

When defining an SMS gateway, three parameters must be specified in order to define the formatting:

- **email_address**: This is the email address used for sending the SMS via the email-to-SMS gateway.
- **subject_line**: This is the subject line of the outgoing email that will be converted to an SMS.
- **email_body**: This is the body of the outgoing email that will be converted to an SMS.

For each of these parameters, the value can be either fixed text, or an event text, or the destination phone number. This is done by embedding the following escape sequences into the text:

- **{areacode}**. This escape sequence is replaced by the destination's cellular number area code.
- **{number}**. This escape sequence is replaced by the destination's cellular local number.
- **{message}**. This escape sequence is replaced by the text to be shown to the user.
- **\{, \}, **. These are replaced by the {, } or \ respectively.

By default, the email to the email-to-SMS server is sent through the defined SMTP servers, prioritized by the Prioritizing SMTP gateways command. If needed, the user may define a specific SMTP gateway or gateways to be used for sending email to this email-to-SMS gateway.

Several SMS gateways can be defined. The system will try the SMS gateways, in the order specified in Prioritizing SMS gateways until it successfully connects to one of them. It is possible to define that specific SMS destinations will use specific SMS gateways (see Defining a new event notification destination).

**Example:**

```
smsgw_define smsgw=SMSGW1 email_address={areacode}{number}@sms2emailserver.example.com
subject_line=SMS email_
body={message}
```

**Output:**

```
Command completed successfully.
```

### Access control

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

### Return codes

- **GATEWAY_MAX_REACHED**

  The maximum allowed number of gateways is already reached.
- **SMSGW_CANNOT_BE_DEFINED_WITHOUT_SMTPGW**

  The SMS gateway cannot be defined if no SMTP gateway is defined.
- **GATEWAY_NAME_DOES_NOT_EXIST**

  The gateway name does not exist.
- **GATEWAY_NAME_APPEARS_TWICE**

  The gateway name appears twice on the list.
- **GATEWAY_NAME_ALREADY_EXISTS**

  The gateway name already exists.

## Deleting an SMS gateway

Use the **smsgw_delete** command to delete an SMS gateway.

```
smsgw_delete smsgw=SMSGatewayName
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **smsgw** | Object name | SMS gateway to be deleted. | Y |

A gateway cannot be deleted if it is part of a notification rule or if it is being used by a destination.

Before deleting an SMS gateway, make sure that all alerting events are cleared.

**Example:**

```
smsgw_delete smsgw=external-SMSGW
```

**Output:**

```
Command completed successfully.
```

### Access control

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

### Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_DELETE_SMS_GATEWAY**

  Are you sure you want to delete SMS gateway *Gateway*?

### Return codes

- **CANNOT_CHANGE_EVENT_CONF_WITH_ALERTING_EVENTS**

  Cannot change an event configuration while there are alerting events.

  **Troubleshooting:** Clear all alerting events before changing an event configuration.
- **GATEWAY_NAME_DOES_NOT_EXIST**

  The gateway name does not exist.
- **GATEWAY_USED_BY_DESTINATION**

  The gateway is used by a destination.

## Listing SMS gateways

Use the **smsgw_list** command to list SMS gateways.

```
smsgw_list [ smsgw=SMSGatewayName ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **smsgw** | Object name | Name of SMS gateway to list. | N | All gateways. |

The command lists all SMS gateways, or a specific one. For each SMS gateway, all of its configuration information is listed.

| Field ID | Field output | Default position |
|---|---|---|
| **name** | Name | 1 |
| **email_address** | Email Address | 2 |
| **gateways** | SMTP Gateways | 3 |
| **subject_line** | Subject Line | N/A |
| **email_body** | Email Body | N/A |
| **priority** | Priority | N/A |

**Example:**

```
smsgw_list
```

**Output:**

```
Name    Email Address                                          SMTP Gateways
SMSGW1  {areacode}{number}@sms2emailserver.example.com         all
SMSGW2  {areacode}{number}@example.com                         all
```

### Access control

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

## Prioritizing SMS gateways

Use the **smsgw_prioritize** command to set the priorities of the SMS gateways for sending SMS messages.

```
smsgw_prioritize order=<gw1[,gw2]...>
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **order** | Object name | List of all SMS gateways ordered by priority. | Y |

SMS messages can be sent to cell phones through one of the email-to-SMS gateways in this list. This command determines the order in which the storage system attempts to use these SMS gateways.

Only one gateway is used and subsequent gateways are only tried if the preceding ones in this priority list return an error.

Specific SMS destinations may define their own SMS gateways to be used when sending SMS to these destinations, regardless of this list.

**Example:**

```
smsgw_prioritize order=SMSGW1,SMSGW2
```

**Output:**

```
Command completed successfully
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |

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

### Return codes

- **CANNOT_CHANGE_EVENT_CONF_WITH_ALERTING_EVENTS**

  Cannot change an event configuration while there are alerting events.

  **Troubleshooting:** Clear all alerting events before changing an event configuration.
- **GATEWAY_NAME_APPEARS_TWICE**

  The gateway name appears twice on the list.
- **GATEWAY_NAME_DOES_NOT_EXIST**

  The gateway name does not exist.
- **GATEWAY_NAME_MISSING_FROM_LIST**

  The gateway name is missing from the list.

## Renaming an SMS gateway

Use the **smsgw_rename** command to rename an SMS gateway.

```
smsgw_rename smsgw=SMSGatewayName new_name=Name
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **smsgw** | Object name | SMS gateway to be renamed. | Y |
| **new_name** | Object name | New name for the SMS gateway. | Y |

Before renaming an SMS gateway, make sure that all alerting events are cleared.

**Example:**

```
smsgw_rename smsgw=SMSGW2 new_name=external-SMSGW
```

**Output:**

```
Command completed successfully
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Disallowed |

| User Category | Permission |
|---|---|
| Read-only users | Disallowed |
| Technicians | Disallowed |

### Return codes

- **CANNOT_CHANGE_EVENT_CONF_WITH_ALERTING_EVENTS**

  Cannot change an event configuration while there are alerting events.

  **Troubleshooting:** Clear all alerting events before changing an event configuration.
- **GATEWAY_NAME_ALREADY_EXISTS**

  The gateway name already exists.
- **GATEWAY_NAME_DOES_NOT_EXIST**

  The gateway name does not exist.

## Updating an SMS gateway

Use the **smsgw_update** command to update an SMS gateway.

```
smsgw_update smsgw=SMSGatewayName [ email_address=email ] [ subject_line=SubjectLineScheme ]
 [ email_body=EmailBodyScheme ] [ smtpgw=<SMTPGW1[,SMTPGW2]...|ALL> ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **smsgw** | Object name | SMS gateway name. | Y | N/A |
| **email_address** | Token String | Format for email address. | N | Leave unchanged. |
| **subject_line** | Token String | Format for subject line. | N | Leave unchanged. |
| **email_body** | Token String | Format for the email's body. | N | Leave unchanged. |
| **smtpgw** | Object name | List of SMTP gateways to be used. | N | The SMTP gateways defined in the smtpgw_prioritize command. |

This command updates the configuration information of an existing SMS gateway. For the exact description and documentation of each parameter, see the documentation of Defining an SMS gateway.

This command cannot be executed while there are uncleared alerting events.

Parameters that are not specified will not be changed.

**Example:**

```
smsgw_update smsgw=SMSGW1 email_address={areacode}{number}@sms2emailserver.example.com
subject_line=NextraSMS
email_body={message}
```

**Output:**

```
Command completed successfully.
```

### Access control

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

### Return codes

- **CANNOT_CHANGE_EVENT_CONF_WITH_ALERTING_EVENTS**

  Cannot change an event configuration while there are alerting events.

  **Troubleshooting:** Clear all alerting events before changing an event configuration.
- **GATEWAY_NAME_APPEARS_TWICE**

  The gateway name appears twice on the list.
- **GATEWAY_NAME_DOES_NOT_EXIST**

  The gateway name does not exist.

## Defining a new SMTP gateway

Use the **smtpgw_define** command to define an SMTP gateway.

```
smtpgw_define smtpgw=SMTPGatewayName address=Address [ from_address=<email|DEFAULT> ]
 [ reply_to_address=<email|DEFAULT> ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **smtpgw** | Object name | SMTP gateway name. | Y | N/A |
| **address** | N/A | SMTP gateway address (IP or DNS name). | Y | N/A |
| **from_address** | N/A | Sender's email address used for outgoing emails sent through this SMTP server. | N | DEFAULT (system-wide sender's address that applies to all servers). |
| **reply_to_address** | N/A | The reply to address used for outgoing emails sent through this SMTP server. | N | DEFAULT (system-wide reply-to address that applies to all servers). |

Several email gateways can be defined to enable notification of events by email or sending SMS messages via email-to-SMS gateways. By default, the system attempts to send each email notification through the first gateway according to the order that you specify. Subsequent gateways are only tried if the first in line returns an error. A specific email destination, or a specific SMS gateway may be defined to use only specific SMTP gateways.

The SMTP protocol dictates that every email message must specify the email address of the sender. This sender address must be a valid address for two reasons:

- Many SMTP gateways require a valid sender address, otherwise they will not forward the email, as a security measure in order to prevent unauthorized usage of the SMTP server. Often this sender address must be limited to a specific domain.
- The sender's address is used as the destination for error messages generated by the SMTP gateways, such as: incorrect email address, full email mailbox and so on.

If the sender's address is not specified for a specific SMTP gateway, a global system-wide sender's address specified in Setting configuration parameters is used.

The user can also configure a reply-to address which is different from the sender's address, if it is required that the return emails be sent to another destination.

**Example:**

```
smtpgw_define smtpgw=mailserver1 address=smtp.example.com from_address=nextra@example.com
reply_to_address=ne
xtraerrors@example.com
```

**Output:**

```
Command completed successfully.
```

### Access control

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

### Return codes

- **FROM_ADDRESS_NOT_DEFINED**

  Neither the gateway's From Address nor the default From Address is defined.
- **GATEWAY_MAX_REACHED**

  The maximum allowed number of gateways is already reached.
- **GATEWAY_NAME_ALREADY_EXISTS**

  The gateway name already exists.

## Deleting an SMTP gateway

Use the **smtpgw_delete** command to delete the specified SMTP gateway.

```
smtpgw_delete smtpgw=SMTPGatewayName
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **smtpgw** | Object name | SMTP gateway to be deleted. | Y |

A gateway cannot be deleted if it is part of a notification rule, is being used as an SMS gateway, or if it belongs to a destination.

Before deleting an SMTP gateway, make sure that all alerting events are cleared.

**Example:**

```
smtpgw_delete smtpgw=mailserverbackup
```

**Output:**

```
Command completed successfully
```

### Access control

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

### Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_DELETE_SMTP_GATEWAY**

  Are you sure you want to delete SMTP gateway *Gateway*?

### Return codes

- **CANNOT_CHANGE_EVENT_CONF_WITH_ALERTING_EVENTS**

  Cannot change an event configuration while there are alerting events.

  **Troubleshooting:** Clear all alerting events before changing an event configuration.
- **GATEWAY_NAME_DOES_NOT_EXIST**

  The gateway name does not exist.
- **GATEWAY_USED_BY_DESTINATION**

  The gateway is used by a destination.
- **GATEWAY_USED_BY_SMS_GATEWAY**

  The gateway is used by an SMS Gateway.

## Listing SMTP gateways

Use the **smtpgw_list** command to list SMTP gateways.

```
smtpgw_list [ smtpgw=SMTPGatewayName ]  [ internal=<yes|no> ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **smtpgw** | Object name | Name of SMTP gateway to list. | N | no. |
| **internal** | Enumeration | Filters gateways by their XIV-internal attribute. | N | no |

This command lists defined SMTP gateways and their configuration information.

| Field ID | Field output | Default position |
|---|---|---|
| **name** | Name | 1 |
| **address** | Address | 2 |
| **priority** | Priority | 3 |
| **from_address** | From Address | N/A |
| **reply_to_address** | Reply-to Address | N/A |
| **failed** | Failed | N/A |
| **port** | Port | N/A |
| **creator** | Creator | N/A |

### Example:

```
smtpgw_list
```

### Output:

```
Name         Email Address          Port  Priority
-----------  --------------------   ----  --------
mailserver1  smtp.example.com       25    1
mailserver2  smtp.example.com       25    2
```

### Access control

| User Category | Permission | Condition |
|---|---|---|
| Storage administrator | Conditionally Allowed | Allowed, unless the internal parameter is specified. |
| Storage integration administrator | Disallowed | N/A |
| Application administrator | Conditionally Allowed | Allowed, unless the internal parameter is specified. |
| Security administrator | Disallowed | N/A |
| Read-only users | Conditionally Allowed | Allowed, unless the internal parameter is specified. |
| Technicians | Allowed | N/A |

## Prioritizing SMTP gateways

Use the **smtpgw_prioritize** command to prirotize SMTP gateways.

```
smtpgw_prioritize order=<gw1[,gw2]...>
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **order** | Object name | List of all the SMTP gateways in the order of their priority. | Y |

Several email gateways can be defined to enable notification of events or the sending of SMS by email. By default, the storage system attempts to send each email through the first gateway according to the order that is specified in this command. Only one gateway is used and subsequent gateways are only tried if the preceding ones in this priority list return an error.

These priorities are used only for email destinations and SMS gateways that did not specify their own SMTP gateways.

**Example:**

```
smtpgw_prioritize order=mailserver2,mailserver1
```

**Output:**

```
Command completed successfully
```

### Access control

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

### Return codes

- **CANNOT_CHANGE_EVENT_CONF_WITH_ALERTING_EVENTS**

  Cannot change an event configuration while there are alerting events.

  **Troubleshooting:** Clear all alerting events before changing an event configuration.
- **GATEWAY_NAME_APPEARS_TWICE**

  The gateway name appears twice on the list.
- **GATEWAY_NAME_DOES_NOT_EXIST**

  The gateway name does not exist.
- **GATEWAY_NAME_MISSING_FROM_LIST**

  The gateway name is missing from the list.

## Renaming an SMTP gateway

Use the **smtpgw_rename** command to rename an SMTP gateway.

```
smtpgw_rename smtpgw=SMTPGatewayName new_name=Name
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **smtpgw** | Object name | SMTP gateway to be renamed. | Y |
| **new_name** | Object name | New name for the SMTP gateway. | Y |

**Example:**

```
smtpgw_rename smtpgw=mailserver2 new_name=mailserverbackup
```

**Output:**

```
Command completed successfully.
```

### Access control

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

### Return codes

- **CANNOT_CHANGE_EVENT_CONF_WITH_ALERTING_EVENTS**

  Cannot change an event configuration while there are alerting events.

  **Troubleshooting:** Clear all alerting events before changing an event configuration.
- **GATEWAY_NAME_ALREADY_EXISTS**

  The gateway name already exists.
- **GATEWAY_NAME_DOES_NOT_EXIST**

  The gateway name does not exist.

## Updating an SMTP gateway

Use the **smtpgw_update** command to update the configuration of an SMTP gateway.

```
smtpgw_update smtpgw=SMTPGatewayName [ address=Address ] [ from_address=<email|DEFAULT> ]
 [ reply_to_address=<email|DEFAULT> ] [ internal=<yes|no> ] [ port=PortNumber ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **smtpgw** | Object name | SMTP gateway name. | Y | N/A |
| **address** | N/A | SMTP gateway address (IP or DNS name). | N | Leave unchanged. |
| **internal** | Boolean | For an XIV internal gateway, set to Yes. | N | NO |

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **from_address** | N/A | Sender's email address used for out-going emails sent through this SMTP server, or DEFAULT for the system-wide default. | N | Leave unchanged. |
| **reply_to_address** | N/A | The reply-to address used for outgoing emails sent through this SMTP server, or DEFAULT for the system-wide default. | N | Leave unchanged. |
| **port** | Integer | TCP port used in the gateway instead of the default port 25. | N | Leave unchanged. |

This command updates the configuration of an existing SMTP gateway. Fields which are not specified are not changed.

**Example:**

```
smtpgw_update smtpgw=mailserver1 address=smtp2.example.com from_address=nextra@example.com
reply_to_address=ne
xtraerrors@example.com
```

**Output:**

```
Command completed successfully.
```

## Access control

| User Category | Permission | Condition |
|---|---|---|
| Storage administrator | Conditionally Allowed | Allowed, unless the internal parameter is specified. |
| Storage integration administrator | Disallowed | N/A |
| Application administrator | Conditionally Allowed | Allowed, unless the internal parameter is specified. |
| Security administrator | Disallowed | N/A |
| Read-only users | Disallowed | N/A |
| Technicians | Allowed | N/A |

## Return codes

- **GATEWAY_NAME_DOES_NOT_EXIST**

  The gateway name does not exist.
- **CANNOT_CHANGE_EVENT_CONF_WITH_ALERTING_EVENTS**

  Cannot change an event configuration while there are alerting events.

  **Troubleshooting:** Clear all alerting events before changing an event configuration.
- **FROM_ADDRESS_NOT_DEFINED**

  Neither the gateway's From Address nor the default From Address is defined.

# Generating an XMPNS admin control event

Use the **xmpns_admin_config_set** command to generate an XMPNS_ADMIN_CONTROL event.

```
xmpns_admin_config_set action=Action user=User
```

## Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **action** | String | Action code text. | Y |
| **user** | String | User name. | Y |

This command generates an **XMPNS_ADMIN_CONTROL** event which includes the `action_codec` text in the event's description field. The username is also added to the action string sent in the description field.

**Example:**

```
xmpns_admin_config_set action user
```

**Output:**

```
Command completed successfully.
```

## Access control

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

# Chapter 15. IP configuration commands

This section describes the command-line interface (CLI) for IP configuration.

## Updating the configuration of the Ethernet port

Use the **ethernet_port_update** command to update the configuration of the Ethernet port.

```
ethernet_port_update ethernet_port_name=EthernetPortName mtu=MTU [ force_mtu_change=<yes|no> ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **ethernet_port_name** | Object name | Name of the Ethernet port. The name format is ethernet_m_p. 'm' is the module_id and 'p' is the port. | Y | N/A |
| **mtu** | Integer | Maximum Transmission Unit: The packet size that is supported by the connecting Ethernet switch. Valid values 1500 - 9000. | Y | N/A |
| **force_mtu_change** | Boolean | If yes, forces the MTU changes for management or VPN Ethernet port. | N | no |

This command updates the configuration of the Ethernet port.

**Example:**

```
ethernet_port_update ethernet_port_name=ethernet_port_4_5 mtu=2000
```

**Output:**

```
Command completed successfully
```

**Attention:**

Be careful when changing Ethernet port MTU. If management connectivity is on, or a VPN is being used, this might break existing application connectivity, unless the adjacent switch is properly configured. This would be the result of the switch using a fixed MTU size lower than the new size being defined on the system.

For example, if the system sends 9000 byte packets while the switch can only receive packets up to 1500 bytes, the switch drops the extra packets.

### Access control

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

© Copyright IBM Corp. 2016, 2022

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

### Return codes

- **ETHERNET_PORT_DOES_NOT_EXIST**

  The Ethernet port does not exist.
- **ETHERNET_PORT_MTU_UPDATE_NOT_SUPPORTED_FOR_MANAGEMENT_OR_VPN**

  Ethernet port MTU update not supported for management or VPN ports.

  **Troubleshooting:** Changing Ethernet port MTU with management or VPN on, might break existing applications connectivity. Verify that the adjacent switch supports the newly configured MTU size.

## Showing the status and configuration of Ethernet ports (deprecated)

Use the **ipinterface_list_ports** command to list all Ethernet ports together with their configurations and statuses.

```
ipinterface_list_ports
```

**Note:**

As of version 12.3.2, this is a legacy command. Instead, use the `ethernet_port_list` command.

All physical Ethernet ports used to connect to the user's network are listed. The list includes the following information:

- Component ID (Module number for iSCSI or switch number for management/field technician port)
- Port number on module/switch
- For management/VPN/field technician: "management"/"VPN"/"field technician"
- IP interface containing the ports (or none, if port is not configured as part of IP interface)
- Status up/down
- Auto-negotiation: Half-full duplex, 1000/100/10

**Example:**

```
ipinterface_list_ports
```

**Output:**

```
Index   Role         IP Interface   Connected Component    Link Up?
------- ------------ -------------- ---------------------- ----------
1       Component                   1:Flash_Canister:4:1   yes
1       Component                   1:Flash_Canister:4:2   yes
1       IPMI                        1:Module:13            yes
1       IPMI                        1:Module:14            yes
1       IPMI                        1:Module:9             yes
1       Internal                    1:IB_Switch:1:12       yes
1       Internal                    1:IB_Switch:1:13       yes
1       Internal                    1:IB_Switch:1:8        yes
1       Management                                         yes
1       iSCSI                                              unknown
1       iSCSI                                              unknown
1       iSCSI                                              unknown
2       IPMI                        1:Module:11            yes
2       IPMI                        1:Module:12            yes
2       IPMI                        1:Module:7             yes
2       iSCSI                                              unknown
2       iSCSI                                              unknown
2       iSCSI                                              unknown

Cont.:

Negotiated Speed (Mb/s)   Full Duplex?   Module        RX Flow Control?   TX Flow Control?
------------------------- -------------- ------------- ------------------ ------------------
1000                      yes            1:Module:12   yes                yes
1000                      yes            1:Module:13   yes                yes
1000                      yes            1:Module:12   yes                yes
1000                      yes            1:Module:13   yes                yes
1000                      yes            1:Module:8    yes                yes
10000                     yes            1:Module:12   yes                yes
10000                     yes            1:Module:13   yes                yes
10000                     yes            1:Module:8    yes                yes
1000                      yes            1:Module:12   yes                yes
N/A                       unknown        1:Module:12   yes                yes
N/A                       unknown        1:Module:13   yes                yes
N/A                       unknown        1:Module:8    yes                yes
1000                      yes            1:Module:12   yes                yes
1000                      yes            1:Module:13   yes                yes
1000                      yes            1:Module:8    yes                yes
N/A                       unknown        1:Module:12   yes                yes
N/A                       unknown        1:Module:13   yes                yes
N/A                       unknown        1:Module:8    yes                yes
```

| Field ID | Field output | Default position |
|---|---|---|
| **index** | Index | 1 |
| **role** | Role | 2 |
| **ip_interface_name** | IP Interface | 3 |
| **connected_component** | Connected Component | 4 |
| **is_link_up** | Link Up? | 5 |
| **negotiated_speed_Mbs** | Negotiated Speed (Mb/s) | 6 |
| **is_full_duplex** | Full Duplex? | 7 |
| **module_id** | Module | 8 |
| **requires_service** | Requires Service | N/A |
| **service_reason** | Service Reason | N/A |
| **pause_autonegotiate** | Flow control auto-negotiate? | N/A |
| **pause_rx** | RX Flow Control? | 9 |
| **pause_tx** | TX Flow Control? | 10 |

## Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |

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

## Showing the status and configuration of Ethernet ports

Use the **ethernet_port_list** command to list all Ethernet ports together with their configurations and statuses.

```
ethernet_port_list
```

All physical Ethernet ports used to connect to the user's network are listed. The list includes the following information:

- Component ID (Module number for iSCSI or switch number for management/field technician port)
- Port number on module/switch
- For management/VPN/field technician: "management"/"VPN"/"field technician"
- IP interface containing the ports (or none, if port is not configured as part of IP interface)
- Status up/down
- Auto-negotiation: Half-full duplex, 1000/100/10

**Example:**

```
ethernet_port_list
```

**Output:**

```
Port Name           Role                  Connected Component    Link Up?
------------------- --------------------- ---------------------- ----------
ethernet_port_1_1   IPMI_1                1:Module:2             yes
ethernet_port_1_2   IPMI_2                1:Module:3             yes
ethernet_port_1_3   Management                                   yes
ethernet_port_1_4   canister_eth_port_2   1:Flash_Canister:1:2   yes
ethernet_port_1_5   Data                                         yes
ethernet_port_1_6   Data                                         no
ethernet_port_2_1   IPMI_1                1:Module:3             yes
ethernet_port_2_2   IPMI_2                1:Module:1             yes
ethernet_port_2_3   Management                                   yes
ethernet_port_2_4   canister_eth_port_1   1:Flash_Canister:1:1   yes
ethernet_port_2_5   Data                                         yes
ethernet_port_2_6   Data                                         no
ethernet_port_3_1   IPMI_1                1:Module:1             yes
ethernet_port_3_2   IPMI_2                1:Module:2             yes
ethernet_port_3_3   Management                                   yes
ethernet_port_3_5   Data                                         yes
ethernet_port_3_6   Data                                         no

cont:

Negotiated Speed (Mb/s)     MTU    Full Duplex?   Module
-------------------------   ------ -------------- ------------
1000                        1500   yes            1:Module:1
1000                        1500   yes            1:Module:1
1000                        1500   yes            1:Module:1
1000                        1500   yes            1:Module:1
10000                       1500   yes            1:Module:1
0                           1500   no             1:Module:1
1000                        1500   yes            1:Module:2
1000                        1500   yes            1:Module:2
1000                        1500   yes            1:Module:2
1000                        1500   yes            1:Module:2
10000                       1500   yes            1:Module:2
0                           1500   no             1:Module:2
1000                        1500   yes            1:Module:3
1000                        1500   yes            1:Module:3
1000                        1500   yes            1:Module:3
10000                       1500   yes            1:Module:3
0                           1500   no             1:Module:3

cont:
Flow control auto-negotiate?   RX Flow Control?   TX Flow Control?
------------------------------ ------------------ ------------------
yes                            yes                yes
yes                            yes                yes
yes                            yes                yes
yes                            yes                yes
no                             yes                yes
no                             yes                yes
yes                            yes                yes
yes                            yes                yes
yes                            yes                yes
yes                            yes                yes
no                             yes                yes
no                             yes                yes
yes                            yes                yes
yes                            yes                yes
yes                            yes                yes
no                             yes                yes
no                             yes                yes
```

| Field ID | Field output | Default position |
|---|---|---|
| **ethernet_port_name** | Port Name | 1 |
| **role** | Role | 2 |
| **connected_component** | Connected Component | 3 |
| **is_link_up** | Link Up? | 4 |
| **negotiated_speed_Mbs** | Negotiated Speed (Mb/s) | 5 |
| **mtu** | MTU | 6 |
| **is_full_duplex** | Full Duplex? | 7 |

| Field ID | Field output | Default position |
|---|---|---|
| **module_id** | Module | 8 |
| **requires_service** | Requires Service | N/A |
| **service_reason** | Service Reason | N/A |
| **pause_autonegotiate** | Flow control auto-negotiate? | 9 |
| **pause_rx** | RX Flow Control? | 10 |
| **pause_tx** | TX Flow Control? | 11 |
| **nic_component_id** | NIC | N/A |
| **cna_component_id** | CNA | N/A |

### Access control

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

## Displaying Ethernet (physical) port statistics in the system

Use the **ethernet_port_counter_list** command to display Ethernet (physical) port statistics in the system.

```
ethernet_port_counter_list
```

**Example:**

```
ethernet_port_counter_list
```

**Output:**

```
Port Name           Role         Module       Connected Component
------------------- ------------ ------------ ---------------------
ethernet_port_1_1   IPMI         1:Module:1   1:Module:2
ethernet_port_1_2   IPMI         1:Module:1   1:Module:3
ethernet_port_1_3   Management   1:Module:1
ethernet_port_1_4   Component    1:Module:1   1:Flash_Canister:1:2
ethernet_port_1_5   iSCSI        1:Module:1
ethernet_port_1_6   iSCSI        1:Module:1
ethernet_port_2_1   IPMI         1:Module:2   1:Module:3
ethernet_port_2_2   IPMI         1:Module:2   1:Module:1
ethernet_port_2_3   Management   1:Module:2
ethernet_port_2_4   Component    1:Module:2   1:Flash_Canister:1:1
ethernet_port_2_5   iSCSI        1:Module:2
ethernet_port_3_1   IPMI         1:Module:3   1:Module:1
ethernet_port_3_2   IPMI         1:Module:3   1:Module:2
ethernet_port_3_5   iSCSI        1:Module:3
ibtun0              Internal     1:Module:1   1:Module:2
ibtun0              Internal     1:Module:2   1:Module:3
ibtun0              Internal     1:Module:3   1:Module:1

cont:
RX bytes     RX packets   TX bytes    TX packets   RX errors   RX drops
------------ ------------ ----------- ------------ ----------- ----------
814582       8131         814582      8131         0           0
815590       8145         815860      8148         0           0
21626657     271999       2613884     5554         0           0
10044406     18337        3439149     23587        0           0
0            0            0           0            0           0
0            0            0           0            0           0
811622       8083         812154      8089         0           0
814916       8134         815092      8136         0           0
20182785     266284       29959       312          0           0
142025271    106676       7628675     92402        0           0
0            0            0           0            0           0
816288       8152         816370      8153         0           0
812154       8089         811892      8086         0           0
0            0            0           0            0           0
210368770    1587299      537221876   1287096      0           0
246042184    1898524      594705380   1026429      0           0
1099820788   2144701      260338771   1159386      0           0
```

| Field ID | Field output | Default position |
|---|---|---|
| **ethernet_port_name** | Port Name | 1 |
| **role** | Role | 2 |
| **connected_component** | Connected Component | 3 |
| **rx_bytes** | RX bytes | 4 |
| **rx_pkts** | RX packets | 5 |
| **tx_bytes** | TX bytes | 6 |
| **tx_pkts** | TX packets | 7 |
| **rx_errors** | RX errors | 8 |
| **rx_drop** | RX drops | 9 |
| **missing** | missing | N/A |

## Access control

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

| User Category | Permission |
|---|---|
| Technicians | Allowed |

# Creating a new IP interface

Use the **ipinterface_create** command to create a new IP interface for iSCSI ports.

```
ipinterface_create ipinterface=IPInterfaceName address=Address netmask=NetworkMask
[ gateway=DefaultGateway ]
 < < ethernet_port_name=EthernetPortName [ vlan_id=VlanID ]
 > | < module=ModuleNumber port=PortNumber [ mtu=MTU ]
 [ speed=<auto|10mb|100mb|1000mb|1gb|2500mb|2.5gb|10000mb|10gb> ] > >
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **ipinterface** | Object name | The name of the IP interface to be created. Do not use the names `Management` or VPN. | Y | N/A |
| **address** | N/A | IP address of the interface. | Y | N/A |
| **netmask** | N/A | Network mask of the interface. | Y | N/A |
| **gateway** | N/A | IP address of the default gateway for this interface. | N | None |
| **ethernet_port_name** | Object name | Name of the Ethernet port. The name format is ethernet_m_p. 'm' is the module_id and 'p' is the port. | N | N/A |
| **vlan_id** | Integer | VLAN id. Valid values 1-4094. | N | 0 |
| **module** | N/A | Component identifier (rack and module) of the module containing Ethernet ports. | N | N/A |
| **port** | Integer | Port Number. | N | N/A |
| **mtu** | Integer | Maximum Transmission Unit: The supported packet size by the connecting Ethernet switch.<br>This is optional when the default equals 1536. MTU of up to 4500 is supported. | N | 4500 for iSCSI and 1536 for Management and VPN. |
| **speed** | Enumeration | Interface's speed. An explicit speed turns off auto-negotiation. Valid values: 'auto' or 'explicit'. | N | auto |

This command defines a new IP interface for iSCSI traffic. Gateway, MTU, network mask and IP are the standard IP definitions.

Each iSCSI Ethernet port can be defined as an IP interface.

## Access control

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

## Return codes

- **IPINTERFACE_EXISTS**

  This IP interface name is already in use.
- **ILLEGAL_PORT_NUMBER**

  The port number is out of range.
- **ETHERNET_PORT_DOES_NOT_EXIST**

  The Ethernet port does not exist.
- **PORT_IS_USED_IN_ANOTHER_IP_INTERFACE**

  One of the physical ports specified is already assigned to an IP Interface.
- **IP_ADDRESS_ALREADY_USED_IN_ANOTHER_INTERFACE**

  The IP address is already assigned to another interface.
- **IPADDRESS_AND_GATEWAY_ARE_NOT_ON_SAME_SUBNET**

  The IP address specified for the default gateway is not in the subnet of the IP interface.
- **MTU_TOO_LARGE**

  The specified MTU value is too large.
- **ILLEGAL_COMPONENT_ID**

  This component ID is illegal.
- **ILLEGAL_IPADDRESS**

  An illegal IP address was entered.
- **DUPLICATE_IPADDRESSES**

  Duplicate IP addresses were specified.
- **ILLEGAL_GATEWAY_IPADDRESS**

  An illegal IP address was specified for the default gateway.
- **TOO_MANY_VLANS**

  The number of VLAN IDs is larger than that allowed by the system.
- **TOO_MANY_IPINTERFACES**

  The number of IP interfaces is too large.
- **TOO_MANY_IPINTERFACES_PER_ETHERNET_PORT**

  The number of IP interfaces per Ethernet port is larger than that allowed by the system.
- **TOO_MANY_IPINTERFACES_PER_VLAN**

  The number of IP interfaces per VLAN is larger than that allowed by the system.
- **IP_ADDRESS_ALREADY_EXISTS**

The system already has the same IP address on the same VLAN.

- **ILLEGAL_ETHERNET_PORT_WITH_VLAN**

  The Ethernet port role does not support VLANs. Only management, VPN, or iSCSI roles are allowed.
- **TWO_IPINTERFACES_ON_SAME_ETHERNET_PORT_AND_VLAN_NOT_ALLOWED**

  Two IP interfaces on the same Ethernet port and having the same VLAN ID are not allowed.
- **MAX_IP_INTERFACES_FOR_TARGET_CONNECTIVITY_PER_ETHERNET_PORT_REACHED**

  Maximal number of IP Interfaces for target connectivity using same port reached.
- **PARAMETER_CANNOT_BE_UPDATED_ON_IPINTERFACE_WITH_VLAN**

  The MTU, speed, module or port parameter can be used only if VLAN is not configured and this is the only IP interface on the physical port.

# Deleting IP interfaces

Use the **ipinterface_delete** command to delete an IP interface.

```
ipinterface_delete ipinterface=IPInterfaceName
```

## Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **ipinterface** | Object name | The IP interface to be deleted. | Y |

Only the interfaces defined for iSCSI traffic can be deleted. Management and VPN interfaces cannot be deleted.

## Access control

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

## Return codes

- **IPINTERFACE_DOES_NOT_EXIST**

  This IP interface name does not exist.
- **COMMAND_NOT_ALLOWED_ON_MANAGEMENT_OR_VPN_INTERFACE**

  The operation is not allowed on the management or VPN IP Interface.
- **IPINTERFACE_HAS_CONNECTIVITY**

  The IP interface has connectivity defined to another machine.

# Listing IP interface configuration

Use the **ipinterface_list** command to list the configuration of a specific IP interface or all IP interfaces.

```
ipinterface_list [ ipinterface=IPInterfaceName | address=Address | address6=IPv6address ]
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **ipinterface** | Object name | The IP interface to be listed. | N | All interfaces |
| **address** | N/A | IP address of the interface to be listed. | N | All interfaces |
| **address6** | N/A | IPv6 address of the interface to be listed. | N | All interfaces |

This command lists configuration information for the specified IP interface, or for all IP interfaces (including management). The management or VPN name can only be used to view the configuration of the management of VPN interfaces.

The following information is listed:

- Name
- Type (iSCSI/management)
- IP address (or comma separated addresses for management and VPN)
- Network mask
- Default gateway
- CIDR address (or comma separated addresses for management and VPN)
- Default IPv6 gateway
- MTU
- Module (for iSCSI only)
- Comma separated list of ports (for iSCSI only)
- Interface desired speed information

**Example:**

```
ipinterface_list
```

**Output:**

```
Name         Type         IP Address      Network Mask    Default Gateway   IPv6 Address
------------ ------------ --------------- --------------- ----------------- --------------
management   Management   192.0.2.1       255.255.255.0   192.0.2.254

Cont.:

IPv6 Gateway   MTU    Module        Port   IP access group name
-------------- ------ ------------- ------ ---------------------
               1500   1:Module:12
```

| Field ID | Field output | Default position |
|---|---|---|
| **name** | Name | 1 |
| **ethernet_port_name** | Physical Port Name | 2 |

| Field ID | Field output | Default position |
|---|---|---|
| **type** | Type | 3 |
| **vlan_id** | VLAN ID | 4 |
| **vlan_pcp** | VLAN PCP | 5 |
| **address** | IP Address | 6 |
| **netmask** | Network Mask | 7 |
| **gateway** | Default Gateway | 8 |
| **address6** | IPv6 Address | 9 |
| **gateway6** | IPv6 Gateway | 10 |
| **mtu** | MTU | 11 |
| **module** | Module | 12 |
| **port** | Port | 13 |
| **speed** | Speed | N/A |
| **access_group** | IP access group name | 14 |

### Access control

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

## Listing IP interface addresses

Use the **ipinterface_list_ips** command to list the IP addresses configured on a specific IP interface or all IP interfaces.

```
ipinterface_list_ips [ ipinterface=IPInterfaceName | address=Address | address6=IPv6address |
module=ModuleNum
ber ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **ipinterface** | Object name | The IP interface to be listed. | N | All interfaces |
| **address** | N/A | IP address of the interface to be listed. | N | All addresses |
| **address6** | N/A | IPv6 address of the interface to be listed. | N | All addresses |
| **module** | N/A | Limits the listing to a specific module. | N | All modules |

This command lists IP addresses for the specified interface, or for the specified module, or for both (including Management). The Management or VPN name can only be used to view IP addresses configured for the management of VPN interfaces.

The following information is listed:

- IP Interface
- Interface Type (iSCSI/Management/VPN)
- Address (in CIDR format)
- Address type (Static IPv4/Static IPv6/Link Local IPv6/Site Local IPv6/Global IPv6)
- Module

**Example:**

```
ipinterface_list_ips
```

**Output:**

```
IP Interface   Interface Type   Address
-------------- ---------------- ----------------------------------------
management     Management       2001:0DB8::0001/32
management     Management       192.0.2.1/24
management     Management       2001:0DB8:42f2:e9ff:feaf:ccb2/32

Cont.:

Address Type      Module        IP access group name
----------------- ------------- ---------------------
Global IPv6       1:Module:12
Static IPv4       1:Module:12
Link Local IPv6   1:Module:12
```

| Field ID | Field output | Default position |
|---|---|---|
| **ipinterface** | IP Interface | 1 |
| **ipinterface_type** | Interface Type | 2 |
| **address** | Address | 3 |
| **address_type** | Address Type | 4 |
| **module** | Module | 5 |
| **access_group** | IP access group name | 6 |
| **vlan_id** | VLAN ID | 7 |

### Access control

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

## Displaying IP interface statistics in the system

Use the **ipinterface_counter_list** command to display the IP interfaces and their statistics in the system.

```
ipinterface_counter_list
```

**Example:**

```
ipinterface_counter_list
```

**Output:**

```
Role         Name                   RX packets   TX packets   RX errors
------------ ---------------------- ------------ ------------ -----------
Management   ethernet_port_1_3      244553       5580         0
Management   ethernet_port_2_3      239749       324          0
Management   ethernet_port_3_3      239768       328          0
iSCSI        test_iscsi_m2_p2       0            0            0
iSCSI        test_old_iscsi_m2_p2   0            0            0
```

| Field ID | Field output | Default position |
|---|---|---|
| **role** | Role | 1 |
| **name** | Name | 2 |
| **rx_pkts** | RX packets | 3 |
| **tx_pkts** | TX packets | 4 |
| **rx_errors** | RX errors | 5 |

### Access control

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

## Renaming an IP interface

Use the **ipinterface_rename** command to rename an IP interface.

```
ipinterface_rename ipinterface=IPInterfaceName new_name=Name
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **ipinterface** | Object name | Original name of the IP interface. | Y |
| **new_name** | Object name | The new name of the IP interface. | Y |

This command renames an IP interface. The IP interface must be unique in the system. This command cannot be applied to Management or VPN interfaces.

### Access control

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

### Return codes

- **IPINTERFACE_DOES_NOT_EXIST**

  This IP interface name does not exist.
- **IPINTERFACE_EXISTS**

  This IP interface name is already in use.
- **COMMAND_NOT_ALLOWED_ON_MANAGEMENT_OR_VPN_INTERFACE**

  The operation is not allowed on the management or VPN IP Interface.

## Printing the ARP database of an IP interface

Use the **ipinterface_run_arp** command to print the ARP database of the specified IP interface.

```
ipinterface_run_arp < localipaddress=IPaddress [ vlan_id=VlanID ] > | localipaddress6=IPv6address
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **localipaddress** | N/A | IP address of the IP interface for the ARP database to be printed. | N | N/A |
| **vlan_id** | Integer | The virtual LAN ID. Valid values 1-4094. | N | 0 |
| **localipaddress6** | N/A | IPv6 address of the IP interface for the ARP database to be printed. | N | N/A |

This command prints a list of the ARP database of an IP interface with its IP addresses and their associated Ethernet MAC addresses. The IP address must be one of the IP addresses defined for iSCSI IP interfaces, or the management or VPN name.

| Field ID | Field output | Default position |
|---|---|---|
| **arp_output** | arp Output | 1 |

### Access control

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

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

### Return codes

- **NO_IP_INTERFACE_MATCHES_CRITERIA**

  No IP Interface or VLAN ID matches the defined criteria. Check if the IP Address and/or VLAN ID are correct.
- **MORE_THAN_ONE_IP_INTERFACE_MATCHES**

  More than one IP Interface matches the defined criteria.

## Testing the traceroute to a remote IP

Use the **ipinterface_run_traceroute** to test connectivity to a remote IP node using the ICMP trace-route mechanism.

```
ipinterface_run_traceroute localipaddress=IPaddress remote=remoteHost [ vlan_id=VlanID ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **localipaddress** | N/A | IP address of the IP interface for which the traceroute command is run. | Y | N/A |
| **remote** | N/A | IP address or DNS for the traceroute test. | Y | N/A |
| **vlan_id** | Integer | The virtual LAN ID. Valid values 1-4094. | N | 0 |

This command runs a route trace to the specified remote host through the specified IP interface. The IP address must be one of the IP addresses defined for iSCSI IP interfaces or the Management or VPN name.

| Field ID | Field output | Default position |
|---|---|---|
| **traceroute_output** | traceroute Output | 1 |

### Access control

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

### Return codes

- **NO_IP_INTERFACE_MATCHES_CRITERIA**

No IP Interface or VLAN ID matches the defined criteria. Check if the IP Address and/or VLAN ID are correct.

- **MORE_THAN_ONE_IP_INTERFACE_MATCHES**

  More than one IP Interface matches the defined criteria.

## Testing the traceroute to a remote IPv6 address

Use the **ipinterface_run_traceroute6** command to test connectivity to a remote IPv6 address using the ICMP trace-route mechanism.

```
ipinterface_run_traceroute6 localipaddress6=IPv6address remote6=remoteHost
```

### Parameters

| Name | Description | Mandatory |
|---|---|---|
| **localipaddress6** | IPv6 address of the IP interface for which the traceroute6 command is run. | Y |
| **remote6** | IPv6 address or DNS for the traceroute test. | Y |

This command runs a route trace to the specified remote host through the specified IP interface. The IP address must be one of the IP addresses defined for iSCSI IP interfaces or the Management or VPN name.

| Field ID | Field output | Default position |
|---|---|---|
| **traceroute_output** | traceroute Output | 1 |

### Access control

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

### Return codes

- **NO_IP_INTERFACE_MATCHES_CRITERIA**

  No IP Interface or VLAN ID matches the defined criteria. Check if the IP Address and/or VLAN ID are correct.

- **MORE_THAN_ONE_IP_INTERFACE_MATCHES**

  More than one IP Interface matches the defined criteria.

## Updating an IP interface

Use the **ipinterface_update** command to update the configuration of an IP interface.

```
ipinterface_update ipinterface=IPInterfaceName [ address=Address ] [ netmask=NetworkMask ]
 [ gateway=DefaultGateway ] [ address6=IPv6address ] [ gateway6=DefaultIPv6Gateway ]
 [ access_group=IPAccessGroupName ] [ mtu=MTU ] [ vlan_id=VlanID ]
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **ipinterface** | Object name | The name of the IP interface to be updated. | Y | N/A |
| **address** | N/A | IP address of the interface or a list of addresses for the Management and VPN interfaces. | N | Leaves the address unchanged. |
| **netmask** | N/A | Network mask of the interface. | N | Leaves the network mask unchanged. |
| **gateway** | N/A | IP address of the default gateway for this interface. | N | Leaves unchanged. |
| **address6** | N/A | IPv6 address of the interface or a list of addresses for the Management and VPN interfaces. | N | Leaves the address unchanged. |
| **gateway6** | N/A | IPv6 address of the default gateway for this interface. | N | Leaves unchanged. |
| **mtu** | Integer | Maximum Transmission Unit: The packet size that is supported by the connecting Ethernet switch. Valid values 1500 - 9000. | N | Keep unchanged. |
| **access_group** | Object name | The name of the IP access group used for IP filtering. | N | Keep unchanged. |
| **vlan_id** | Integer | The virtual LAN ID. Valid values 1-4094. | N | 0 |

This command updates the configuration of an existing IP interface.

Fields that are not specified do not change their values.

The name of the interface may either be one of the previously defined IP interfaces for iSCSI, or Management for the management IP interface, or VPN for the VPN interface.

Management ports are dedicated for CLI and GUI communications, as well as for outgoing SNMP and SMTP connections. For management interfaces, the user must specify three IP addresses (equal to the number of potential managers, minus the number of management ports).

For VPN interfaces, the user must specify two IP addresses (equal to the number of VPN ports). All VPN addresses must reside on the same subnet.

The mtu parameter can only be used if VLAN is not enabled and this IP interface is the only one configured on the Ethernet port.

**Example:**

```
ipinterface_update ipinterface=management
```

**Output:**

```
Command completed successfully
```

**Attention:**

Be careful when changing Ethernet port MTU. If management connectivity is on, or a VPN is being used, this might break existing application connectivity, unless the adjacent switch is properly configured. This would be the result of the switch using a fixed MTU size lower than the new size being defined on the system.

For example, if the system sends 9000 byte packets while the switch can only receive packets up to 1500 bytes, the switch drops the extra packets.

## Access control

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

## Return codes

- **IPINTERFACE_DOES_NOT_EXIST**

  This IP interface name does not exist.
- **IP_ADDRESS_ALREADY_USED_IN_ANOTHER_INTERFACE**

  The IP address is already assigned to another interface.
- **IPADDRESS_AND_GATEWAY_ARE_NOT_ON_SAME_SUBNET**

  The IP address specified for the default gateway is not in the subnet of the IP interface.
- **IPINTERFACE_MANAGEMENT_DIFFERENT_SUBNET**

  All IP addresses management modules must be in the same subnet.
- **IPINTERFACE_MANAGEMENT_MISSING_IPS**

  The number of IP addresses specified is smaller than the number of management modules.
- **IPINTERFACE_MANAGEMENT_TOO_MANY_IPS**

  The number of IP addresses specified is larger than the number of management modules.
- **MTU_TOO_LARGE**

  The specified MTU value is too large.
- **ILLEGAL_IPADDRESS**

  An illegal IP address was entered.
- **DUPLICATE_IPADDRESSES**

  Duplicate IP addresses were specified.
- **ILLEGAL_GATEWAY_IPADDRESS**

  An illegal IP address was specified for the default gateway.
- **ILLEGAL_IPV6ADDRESS**

  An illegal IPv6 address was entered.
- **DUPLICATE_IPV6ADDRESSES**

  Duplicate IPv6 addresses were specified.
- **ILLEGAL_GATEWAY_IPV6_ADDRESS**

An illegal IPv6 address was specified for the default gateway.

- **IPV6ADDRESS_AND_GATEWAY_ARE_NOT_ON_SAME_SUBNET**

  The IPv6 address specified for the default gateway is not in the subnet of the IP interface.
- **IPV6_ADDRESS_ALREADY_USED_IN_ANOTHER_INTERFACE**

  The IPv6 address is already assigned to another interface.
- **IPINTERFACE_MANAGEMENT_MISSING_IPV6S**

  The number of IPv6 addresses specified is smaller than the number of management modules.
- **IPINTERFACE_MANAGEMENT_TOO_MANY_IPV6S**

  The number of IPv6 addresses specified is larger than the number of management modules.
- **IPINTERFACE_MANAGEMENT_DIFFERENT_IPV6_SUBNET**

  All IPv6 addresses management modules must be in the same subnet.
- **IP_ACCESS_GROUP_DOES_NOT_EXIST**

  An IP access group with the specified name does not exist.
- **IP_ACCESS_INVALID_INTERFACE_TYPE**

  IP filtering is applied to an invalid interface (should be management or VPN).
- **IPINTERFACE_VLAN_UPDATE_LIMITED_TO_MANAGEMENT_OR_VPN**

  The VLAN ID can be updated only for management or VPN IP interfaces.
- **PARAMETER_CANNOT_BE_UPDATED_ON_IPINTERFACE_WITH_VLAN**

  The MTU, speed, module or port parameter can be used only if VLAN is not configured and this is the only IP interface on the physical port.
- **ETHERNET_PORT_DOES_NOT_EXIST**

  The Ethernet port does not exist.

# Updating VLAN priority code points

Use the **ipinterface_vlan_update** command to update the priority code points of all IP interfaces which have a VLAN ID defined.

```
ipinterface_vlan_update vlan_pcp=VlanPcp
```

## Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **vlan_pcp** | Integer | The VLAN priority code point (PCP). Valid values 0-7. | Y |

This command updates the priority code points for all IP interfaces which have a VLAN ID defined.

**Example:**

```
ipinterface_vlan_update vlan_pcp=2
```

**Output:**

```
Command completed successfully
```

### Access control

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

## Defining a new IPSec connection

Use the **ipsec_connection_add** command to add a new IPSec connection.

```
ipsec_connection_add ipsec_connection=ConnectionName left=IPInterfaceName
[ right_ip=RightIpAddress ]
 < passkey=PassKey | certificate=PemCertificate >
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **ipsec_connection** | N/A | The name of the IPSec connection to be added. | Y | N/A |
| **left** | Object name | The name of the IP interface to be used as the left side: `management` or VPN. | Y | N/A |
| **right_ip** | N/A | IP address of the right side. | N | Any |
| **passkey** | N/A | Secret password. | N | N/A |
| **certificate** | N/A | The public certificate or certificate chain of the new IPSec connection (see below for details). | N | N/A |

This command defines a new IPSec connection between an IP interface and the right side.

IP interface can be either `management` or VPN. If specified:

- The address of the right side is IPv4 or IPv6; otherwise the right side can be any.
- The secret password must be shared between the left and the right sides.
- The certificate must contain a public key of the right side.

**The certificate parameter:**

The value of the certificate parameter is the content of a PEM file with asterisks instead of newlines. Chained certificates are supported. The total maximal length of a PEM file holding chained certificates (leaf first, root last) is 15360 characters (including the asterisk characters). In Windows, you can drag-and-drop a PEM file from the Windows Explorer to the appropriate location in the XCLI session window; the content will be added automatically.

**Example:**

```
ldap_update_server fqdn=ldap.example.com address=192.0.2.1
        remove_certificate=yes
```

**Example:**

```
ipsec_connection_add ipsec_connection=MySec left=management passkey="MyPass123"
```

## Access control

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

## Return codes

- **IPSEC_CONNECTION_EXISTS**

  The IPSec connection already exists.
- **IPSEC_CONNECTION_BETWEEN_ENDPOINTS_EXISTS**

  A connection between these endpoints already exists.
- **LEFT_INTERFACE_NOT_FOUND**

  The specified left side interface was not found.
- **MAX_IPSEC_CONNECTIONS_REACHED**

  The maximum allowed number of configured IPSec connections is already reached.
- **IPSEC_UNSUPPORTED_FOR_ISCSI**

  IPSec is not supported for iSCSI ports.
- **SSL_CERTIFICATE_CHAIN_EMPTY**

  No certificates were found in the input.
- **SSL_CERTIFICATE_HAS_EXPIRED**

  The SSL certificate has expired.
- **SSL_CERTIFICATE_INVALID_FORMAT**

  The SSL certificate format is invalid or corrupted.
- **SSL_CERTIFICATE_ISSUER_NOT_FOUND**

  The SSL certificate issuer was not found in the certificate chain.
- **SSL_CERTIFICATE_NOT_YET_VALID**

  The SSL certificate is not yet valid.
- **SSL_CERTIFICATE_VERIFICATION_FAILED**

  The SSL certificate chain verification failed.
- **SSL_CERTIFICATE_VERIFICATION_INTERNAL_ERROR**

  The SSL certificate verification has failed because of an internal system error.

# Updating an existing IPSec connection

Use the **ipsec_connection_update** command to update an existing IPSec connection.

```
ipsec_connection_update ipsec_connection=ConnectionName [ left=IPInterfaceName ]
[ right_ip=RightIpAddress ]
 [ passkey=PassKey | certificate=PemCertificate ]
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **ipsec_connection** | Object name | The name of the IPSec connection to be updated. | Y | N/A |
| **left** | Object name | The name of the IP interface to be used as left side: `management` or VPN. | N | None |
| **right_ip** | N/A | The IP address of the right side. | N | None |
| **passkey** | N/A | Pre-shared key. | N | None |
| **certificate** | N/A | The public certificate or certificate chain of the IPSec connection to be updated (see below for details). | N | None |

This command updates an existing IPSec connection between an IP interface and the right side.

IP interface can be either `management` or VPN. If specified:

- The address of the right side is IPv4 or IPv6; otherwise the right side can be any.
- The pre-shared key must be shared between the left and the right sides.
- The certificate must contain a public key of the right side.

**The certificate parameter:**

The value of the certificate parameter is the content of a PEM file with asterisks instead of newlines. Chained certificates are supported. The total maximal length of a PEM file holding chained certificates (leaf first, root last) is 15360 characters (including the asterisk characters). In Windows, you can drag-and-drop a PEM file from the Windows Explorer to the appropriate location in the XCLI session window; the content will be added automatically.

**Example:**

```
ipsec_connection_update ipsec_connection=MySec passkey="MyNewPass!@#"
```

## Access control

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

### Return codes

- **IPSEC_CONNECTION_DOES_NOT_EXIST**

  The specified IPSec connection does not exist.
- **IPSEC_CONNECTION_EXISTS**

  The IPSec connection already exists.
- **LEFT_INTERFACE_NOT_FOUND**

  The specified left side interface was not found.
- **IPSEC_UNSUPPORTED_FOR_ISCSI**

  IPSec is not supported for iSCSI ports.
- **SSL_CERTIFICATE_CHAIN_EMPTY**

  No certificates were found in the input.
- **SSL_CERTIFICATE_HAS_EXPIRED**

  The SSL certificate has expired.
- **SSL_CERTIFICATE_INVALID_FORMAT**

  The SSL certificate format is invalid or corrupted.
- **SSL_CERTIFICATE_ISSUER_NOT_FOUND**

  The SSL certificate issuer was not found in the certificate chain.
- **SSL_CERTIFICATE_NOT_YET_VALID**

  The SSL certificate is not yet valid.
- **SSL_CERTIFICATE_VERIFICATION_FAILED**

  The SSL certificate chain verification failed.
- **SSL_CERTIFICATE_VERIFICATION_INTERNAL_ERROR**

  The SSL certificate verification has failed because of an internal system error.

## Removing an existing IPSec connection

Use the **ipsec_connection_remove** command to remove an existing IPSec connection.

```
ipsec_connection_remove ipsec_connection=ConnectionName
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **ipsec_connection** | Object name | The name of the IPSec connection to be updated. | Y |

**Example:**

```
ipsec_connection_remove ipsec_connection=connect1
```

**Output:**

```
Command completed successfully
```

### Access control

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

### Return codes

- **`IPSEC_CONNECTION_DOES_NOT_EXIST`**

  The specified IPSec connection does not exist.

## Listing IPSec connections

Use the **`ipsec_connection_list`** command to list all or specific IPSec connections.

```
ipsec_connection_list [ ipsec_connection=ConnectionName ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **`ipsec_connection`** | Object name | The IPSec connection(s) to be listed. | N | All IPsec connections |

| Field ID | Field output | Default position |
|---|---|---|
| **`name`** | IPSec Connection | 1 |
| **`type`** | Type | 2 |
| **`left`** | Left Interface | 3 |
| **`right_ip`** | Right Address | 4 |

### Access control

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

# Listing IPSec tunnels

Use the **ipsec_list_tunnels** command to list IPSec tunnels.

```
ipsec_list_tunnels [ ipsec_connection=ConnectionName ] [ left=IPInterfaceName ]
[ left_ip=InterfaceIpAddress ]
 [ right_ip=RightIpAddress ] [ module=ComponentId ]
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **ipsec_connection** | Object name | Lists all IPSec tunnels of this IPSec connection. | N | IPSec tunnels of all IPsec connections |
| **left** | Object name | Lists all IPSec tunnels from this interface. | N | IPsec tunnels from any interface |
| **left_ip** | N/A | Lists all IPSec tunnels from this left IP. | N | IPsec tunnels from any left IP |
| **right_ip** | N/A | Lists all IPSec tunnels from this right IP. | N | IPsec tunnels to any right IP |
| **module** | N/A | Limits the listing to a specific module. | N | All modules |

| Field ID | Field output | Default position |
|---|---|---|
| **name** | IPSec Connection | 1 |
| **type** | Type | 2 |
| **status** | Status | 3 |
| **left** | Left Interface | 4 |
| **left_ip** | Left Address | 5 |
| **right_ip** | Right Address | 6 |
| **module** | Module | 7 |

## Access control

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

# Connecting to a support center

Use the **support_center_connect** command to connect to a support center.

```
support_center_connect [ < timeout=Timeout [ idle_timeout=IdleTimeout ] > | always_on=<yes|no> ]
 [ module=ModuleNumber ] [ password=Password ]
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **timeout** | N/A | Specifies the duration of the session. After the duration elapses, the session will be disconnected. Time is specified in hh:mm format. | N | none |
| **idle_timeout** | N/A | Specifies the idle time for the session after which it will be disconnected. Time is specified in hh:mm format. | N | [timeout] |
| **module** | N/A | The module from which the connection to the support center should be initiated | N | [ the module that handled the CLI request ] |
| **password** | String | A password set by the customer, that needs to be submitted by support services, in order to start a remote support session Format: string, must be 6-12 alpha-numeric characters, and is case-insensitive. | N | none |
| **always_on** | Boolean | Enables a constant connection to the support center (rather than an on-demand connection). | N | none |

If the support center is not defined, the command will fail.

To control the duration of the session, use the parameters **timeout** and **idle_disconnect**.

## Access control

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

## Return codes

- **NO_SUPPORT_CENTERS_ARE_DEFINED**

  No support centers are defined.
- **IDLE_TIMEOUT_MUST_BE_LOWER_THAN_TIMEOUT**

  The idle timeout must be smaller than the regular timeout.
- **MODULE_HAS_NO_SUPPORT_CENTER_PORT**

The specified module does not have a port from which the support center can connect.

- **NO_MODULE_WITH_SUPPORT_CENTER_PORT**

  No module has a port from which the support center can connect.
- **REMOTE_SUPPORT_CLIENT_ALREADY_RUNNING**

  The Remote Support Client is already running.
- **REMOTE_SUPPORT_CLIENT_AUTOMATICALLY_CONNECT_IS_RUNNING**

  The Remote Support Client is running in automatically connect mode.

  **Troubleshooting:** Run support_center_disconnect to stop it.

# Defining a support center

Use the **support_center_define** command to define a support center.

```
support_center_define support_center=SupportCenterName address=Address [ port=port ]
[ priority=priority ]
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **support_center** | Object name | The name of the support center server | Y | N/A |
| **address** | N/A | The IP address of the support center server | Y | N/A |
| **port** | Positive integer | The TCP port to connect to on the support center | N | 22 |
| **priority** | N/A | The priority of the support center (support centers with a higher priority will be connected first) | N | 0 |

**Example:**

```
support_center_define support_center=somewhere address=192.0.2.1
```

**Output:**

```
Command completed successfully
```

## Access control

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

### Return codes

- **MAX_SUPPORT_CENTERS_DEFINED**

  The maximum number of defined support centers is already reached.
- **SUPPORT_CENTER_ALREADY_EXISTS**

  The support center with the indicated name already exists.

  **Troubleshooting:** Enter a different name.

## Deleting a support center

Use the **support_center_delete** command to delete a support center.

```
support_center_delete support_center=SupportCenterName
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **support_center** | Object name | The name of the support center to delete. | Y |

Sessions that belong to this support center are disconnected, even if they are open at the time of deletion.

### Access control

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

### Warnings

- **ARE_YOU_SURE_TO_DELETE_THE_SUPPORT_CENTER**

  Are you sure you want to delete the support center?.

### Return codes

- **SUPPORT_CENTER_NOT_DEFINED**

  The support center is not defined.
- **CANNOT_DELETE_WHILE_SUPPORT_CENTER_IS_RUNNING**

  The support center is running. Disconnect it before deleting.

## Disconnecting from a support center

Use the **support_center_disconnect** command to disconnect the storage system from a support center.

```
support_center_disconnect
```

**Example:**

```
support_center_disconnect
```

**Output:**

```
Command completed successfully
```

### Access control

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

### Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_DISCONNECT_BUSY_REMOTE_SUPPORT**

  Are you sure you want to disconnect the busy remote support connection?

### Return codes

- **REMOTE_SUPPORT_CLIENT_NOT_RUNNING**

  The Remote Support Client is not running.

## Listing support centers

Use the **support_center_list** command to list support centers.

```
support_center_list
```

This command displays the following information about all defined support centers:

- Name
- IP Address
- Port
- Priority

| Field ID | Field output | Default position |
|---|---|---|
| **name** | Name | 1 |
| **address** | Address | 2 |
| **port** | Port | 3 |
| **priority** | Priority | 4 |

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |

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

## Listing statuses of all support centers

Use the **support_center_status** command to list information about the status of the defined support centers.

```
support_center_status
```

**Example:**

```
support_center_status
```

**Output:**

```
State           Connected sessions   Timeout (min)   Module   Connected since
--------------- -------------------- --------------- -------- -----------------
no connection   0                     no timeout
Cont.:
Destination   Auto Connect Active   Always On
------------- --------------------- -----------
              no                    no
```

| Field ID | Field output | Default position |
|---|---|---|
| **state** | State | 1 |
| **connected_support_sessions** | Connected sessions | 2 |
| **minutes_to_timeout** | Timeout (min) | 3 |
| **running_from_module** | Module | 4 |
| **start_time** | Connected since | 5 |
| **destination** | Destination | 6 |
| **automatically_connect_mode** | Auto Connect Active | 7 |
| **stop_automatically_connect** | Stop support center automatically connect | N/A |
| **always_on** | Always On | 8 |

### Access control

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

# Configuring the support center connection to enable automatic connect on restart

Use the **support_center_config** command to configure the connection to a support center automatically.

```
support_center_config automatically_connect=<yes|no> [ connect_through_module1=module ]
 [ connect_through_module2=module ] [ connect_through_module3=module ] [ password=Password ]
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **automatically_connect** | Boolean | Enables the automatic connection to the support center. | Y | N/A |
| **connect_through_module1** | N/A | The first module from which to establish a connection to the support center automatically. | N | Module with first management port |
| **connect_through_module2** | N/A | The second module from which to establish a connection to the support center automatically. | N | Module with second management port |
| **connect_through_module3** | N/A | The third module from which to establish a connection to the support center automatically. | N | Module with third management port |
| **password** | String | A password set by the customer, that needs to be submitted by support services, in order to start a remote support session Password format: case-insensitive string of 6-12 alphanumeric characters. | N | none |

**Example:**

```
support_center_config automatically_connect=yes
```

**Output:**

```
Command completed successfully
```

## Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Disallowed |

| User Category | Permission |
|---|---|
| Read-only users | Disallowed |
| Technicians | Allowed |

### Return codes

- **MODULE_HAS_NO_SUPPORT_CENTER_PORT**

  The specified module does not have a port from which the support center can connect.
- **NO_MODULE_WITH_SUPPORT_CENTER_PORT**

  No module has a port from which the support center can connect.

## Listing the configuration of support center automatically connect

Use the **support_center_config_list** command to display the configuration of the automatic connection to a support center.

```
support_center_config_list
```

**Example:**

```
support_center_config_list
```

**Output:**

```
Enable Auto Conn   First Module   Second Module   Third Module
------------------ -------------- --------------- --------------
yes                1              2               -1
```

| Field ID | Field output | Default position |
|---|---|---|
| **enable_auto_conn** | Enable Auto Conn | 1 |
| **module1_id** | First Module | 2 |
| **module2_id** | Second Module | 3 |
| **module3_id** | Third Module | 4 |
| **automatically_connect_reason** | Auto Conn Reason | N/A |

### Access control

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

# Creating a new IP access group

Use the **ip_access_group_create** command to create a new IP access group.

```
ip_access_group_create access_group=IPAccessGroupName
```

## Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **access_group** | Object name | The name of the IP access group to be created. | Y |

The group may contain up to 20 addresses and can be used to limit network access to a management/VPN interface.

**Example:**

```
ip_access_group_create access_group=IPAccessGroup1
```

**Output:**

```
Command completed successfully
```

## Access control

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

## Return codes

- **IP_ACCESS_GROUP_ALREADY_EXISTS**

  An IP access group with the specified name already exists.
- **IP_ACCESS_MAXIMUM_NUMBER_OF_GROUPS_IS_REACHED**

  The maximum number of IP access groups is already reached.

# Removing an address from an IP access group

Use the **ip_access_group_remove_address** command to delete the IP address of an access group.

```
ip_access_group_remove_address access_group=IPAccessGroupName address=Address
```

## Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **access_group** | Object name | The name of the IP access group. | Y |

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **address** | N/A | The address that should be deleted from the IP access group. | Y |

As a prerequisite for completing this command, the IP address must be defined for the group.

**Example:**

```
ip_access_group_remove_address access_group=IPAccessGroup1 address=192.0.2.1
```

**Output:**

```
Command completed successfully
```

### Access control

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

### Return codes

- **IP_ACCESS_GROUP_DOES_NOT_EXIST**

  An IP access group with the specified name does not exist.
- **IP_ACCESS_ADDRESS_IS_NOT_VALID**

  The given address is not valid.
- **IP_ACCESS_ADDRESS_IS_NOT_IN_GROUP**

  The specified address is not in the group.

## Adding a new address to an IP access group

Use the **ip_access_group_add_address** command to add a new IP to an access group.

```
ip_access_group_add_address access_group=IPAccessGroupName address=Address [ netmask=NetworkMask ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **access_group** | Object name | The name of an IP access group. | Y | N/A |
| **address** | N/A | A valid IP4 address or FQDN to be added to the IP access group. | Y | N/A |
| **netmask** | N/A | The network mask for a network address range. | N | Single IP address range (255.255.255.255). |

The address can be an IP4 address with or without a netmask, or a valid host name (FQDN).

**Example:**

```
ip_access_group_add_address access_group=IPAccessGroup1 address=192.0.2.1
```

**Output:**

```
Command completed successfully
```

### Access control

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

### Return codes

- **IP_ACCESS_GROUP_DOES_NOT_EXIST**

  An IP access group with the specified name does not exist.
- **IP_ACCESS_REMOTE_RESOLVE_ADDRESS_CALL_HAS_FAILED**

  The remote call to resolve an address has failed.
- **IP_ACCESS_MAXIMUM_NUMBER_OF_ADDRESSES_IN_GROUP_IS_REACHED**

  The maximum number of addresses in the IP access group is already reached.

## Deleting an existing IP access group

Use the **ip_access_group_delete** command to delete an IP access group.

```
ip_access_group_delete access_group=IPAccessGroupName
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **access_group** | Object name | The name of the IP access group to be deleted. | Y |

**Example:**

```
ip_access_group_delete access_group=DBGroupNew
```

**Output:**

```
Command completed successfully
```

### Access control

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

### Return codes

- **IP_ACCESS_GROUP_DOES_NOT_EXIST**

  An IP access group with the specified name does not exist.
- **IP_ACCESS_GROUP_IN_USE**

  The group is used for IP filtering.

## Renaming an existing IP access group

Use the **ip_access_group_rename** command to rename an existing IP access group.

```
ip_access_group_rename access_group=IPAccessGroupName new_name=Name
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **access_group** | Object name | Name of the IP access group to be renamed. | Y |
| **new_name** | Object name | A new name of the IP access group. | Y |

**Example:**

```
ip_access_group_rename access_group=DBGroup new_name=DBGroupNew
```

**Output:**

```
Command completed successfully
```

### Access control

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

### Return codes

- **IP_ACCESS_GROUP_DOES_NOT_EXIST**

  An IP access group with the specified name does not exist.
- **IP_ACCESS_GROUP_ALREADY_EXISTS**

  An IP access group with the specified name already exists.

## Listing IP access groups

Use the **ip_access_group_list** command to list IP access groups.

```
ip_access_group_list
```

| Field ID | Field output | Default position |
|---|---|---|
| **name** | Group Name | 1 |
| **addresses.0** | Address 1 | N/A |
| **addresses.1** | Address 2 | N/A |
| **addresses.2** | Address 3 | N/A |
| **addresses.3** | Address 4 | N/A |
| **addresses.4** | Address 5 | N/A |
| **addresses.5** | Address 6 | N/A |
| **addresses.6** | Address 7 | N/A |
| **addresses.7** | Address 8 | N/A |
| **addresses.8** | Address 9 | N/A |
| **addresses.9** | Address 10 | N/A |
| **addresses.10** | Address 11 | N/A |
| **addresses.11** | Address 12 | N/A |
| **addresses.12** | Address 13 | N/A |
| **addresses.13** | Address 14 | N/A |
| **addresses.14** | Address 15 | N/A |
| **addresses.15** | Address 16 | N/A |
| **addresses.16** | Address 17 | N/A |
| **addresses.17** | Address 18 | N/A |
| **addresses.18** | Address 19 | N/A |
| **addresses.19** | Address 20 | N/A |

### Access control

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

# Listing IP access group addresses

Use the **ip_access_group_address_list** command to list IP access group addresses.

```
ip_access_group_address_list
```

This command lists IP access groups and address lists for these groups.

**Example:**

```
ip_access_group_address_list
```

**Output:**

```
Group Name       Address
---------------- ----------------
DBGroup          192.0.2.1
IPAccessGroup1   198.51.100.2
```

| Field ID | Field output | Default position |
|---|---|---|
| **access_group** | Group Name | 1 |
| **address** | Address | 2 |

## Access control

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

© Copyright IBM Corp. 2016, 2022

# Generating a certificate signing request

Use the **pki_generate_csr** command to generate a certificate signing request.

```
pki_generate_csr name=Name subject=Subject
```

## Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **name** | String | The certificate's symbolic name. | Y |
| **subject** | N/A | The subject name for the generated certificate request. The argument must be formatted as /type0=value0/type1=value1/type2=... . | Y |

**Example:**

```
pki_generate_csr name subject
```

| Field ID | Field output | Default position |
|---|---|---|
| **csr** | CSR | 1 |

## Access control

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

## Return codes

- **FAILED_CREATING_CERTIFICATE_SIGNING_REQUEST**

  Failed to generate a certificate signing request.

  **Troubleshooting:** Generate a certificate signing request with a correct subject (for example, '/C=US/CN=IBM').
- **CERTIFICATE_NAME_DOES_NOT_EXIST**

  A certificate with the indicated name was not found.

  **Troubleshooting:** Enter a different name.

# Generating a private key and CSR

Use the **pki_generate_private_key_and_csr** command to generate a private key and CSR.

```
pki_generate_private_key_and_csr name=Name subject=Subject [ bits=Bits ]
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **bits** | Integer | The private key size in bits. It can be between 1024 to 4096. | N | 2048 |
| **name** | String | The certificate's symbolic name. | Y | N/A |
| **subject** | N/A | The subject name for the generated certificate request. The argument must be formatted as /type0=value0/ type1=value1/ type2=... . | Y | N/A |

**Example:**

```
pki_generate_private_key_and_csr name="my_cert"
subject="/C=US/CN=IBM" bits=1024
```

| Field ID | Field output | Default position |
|---|---|---|
| **csr** | CSR | 1 |

## Access control

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

## Return codes

- **FAILED_CREATING_CERTIFICATE_SIGNING_REQUEST**

  Failed to generate a certificate signing request.

  **Troubleshooting:** Generate a certificate signing request with a correct subject (for example, '/C=US/CN=IBM').
- **FAILED_CREATING_PRIVATE_KEY**

  Failed to create a private key.
- **CERTIFICATE_NAME_ALREADY_EXIST**

  A certificate with the indicated name already exists.

  **Troubleshooting:** Enter a different name.
- **CERTIFICATE_CONTAINER_FULL**

  Cannot add any more certificates, the maximum number is already reached.

  **Troubleshooting:** Delete a certificate.

# Deleting a PKI content

Use the **pki_remove** command to delete the PKI content.

```
pki_remove name=Name
```

## Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **name** | String | The certificate's symbolic name. | Y |

**Example:**

```
pki_remove name="my_cert"
```

## Access control

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

## Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_DELETE_CERTIFICATE**

  Are you sure you want to delete the certificate?

## Return codes

- **CERTIFICATE_NAME_DOES_NOT_EXIST**

  A certificate with the indicated name was not found.

  **Troubleshooting:** Enter a different name.
- **DEFAULT_CERTIFICATE_CANNOT_BE_DELETED**

  The default certificate cannot be deleted.

# Changing a PKI symbolic name

Use the **pki_rename** command to change a PKI symbolic name.

```
pki_rename name=Name new_name=Name
```

## Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **name** | String | The current symbolic name. | Y |
| **new_name** | String | The new symbolic name. | Y |

**Example:**

```
pki_rename name="current_name" new_name="my_new_name"
```

### Access control

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

### Return codes

- **CERTIFICATE_NAME_ALREADY_EXIST**

  A certificate with the indicated name already exists.

  **Troubleshooting:** Enter a different name.
- **CERTIFICATE_NAME_DOES_NOT_EXIST**

  A certificate with the indicated name was not found.

  **Troubleshooting:** Enter a different name.

## Importing a signed certificate

Use the **pki_set_pem** command to import a signed certificate in PEM format.

```
pki_set_pem certificate=SignedCertificate  [ services=<xcli [ ,cim ] [ ,ipsec ] ... | ALL |
NONE> ]
```

### Parameters

| Name | Description | Mandatory | Default |
|---|---|---|---|
| **services** | A comma-separated list of services that use this certificate. | N | none |
| **certificate** | The public certificate or certificate chain to be imported (see below for details). | Y | N/A |

As a security precaution, use the **pki_show_security** command to view the certificate in plain text, and make sure that the certificate text under *Signature Algorithm* does not include the string *MD5*. This will help you avoid a "transcript collision" attack, that can force a hash-construction downgrade to MD5 and reduce expected security. For the vulnerability summary, see the National Vulnerability Database.

**The certificate parameter**

The value of the **certificate** parameter is the content of a PEM file with asterisks instead of newlines. Chained certificates are supported. The total maximal length of a PEM file holding chained certificates (leaf first, root last) is 15360 characters (including the asterisk characters). In Windows, you can drag-and-drop a PEM file from the Windows Explorer to the appropriate location in the CLI session window; the content will be added automatically.

**Example:**

```
pki_set_pem certificate=validCertificateChain
```

**Output:**

```
Command completed successfully
```

## Access control

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

## Return codes

- **SERVICE_IS_USING_OTHER_CERTIFICATE**

  Service '*services*' is using another certificate.

  **Troubleshooting:** Edit the certificate previously used by the service.
- **PRIVATE_KEY_ALREADY_HAS_OTHER_CERTIFICATE**

  The private key matching this certificate already has another certificate.

  **Troubleshooting:** To replace the certificate, use the pki_update command.
- **CERTIFICATE_KEY_WAS_NOT_FOUND**

  Failed to set the certificate.

  **Troubleshooting:** Make sure the certificate parameters are correct.
- **SSL_CERTIFICATE_CHAIN_EMPTY**

  No certificates were found in the input.
- **SSL_CERTIFICATE_HAS_EXPIRED**

  The SSL certificate has expired.
- **SSL_CERTIFICATE_INVALID_FORMAT**

  The SSL certificate format is invalid or corrupted.
- **SSL_CERTIFICATE_ISSUER_NOT_FOUND**

  The SSL certificate issuer was not found in the certificate chain.
- **SSL_CERTIFICATE_NOT_YET_VALID**

  The SSL certificate is not yet valid.
- **SSL_CERTIFICATE_VERIFICATION_FAILED**

  The SSL certificate chain verification failed.
- **SSL_CERTIFICATE_VERIFICATION_INTERNAL_ERROR**

  The SSL certificate verification has failed because of an internal system error.

# Importing a PKCS#12 certificate

Use the **pki_set_pkcs12** command to import a PKCS#12 certificate.

```
pki_set_pkcs12 name=Name password=Password certificate=Base64Data [ services=<xcli [ ,cim ]
[ ,ipsec ]
 ... | ALL | NONE> ]
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **services** | N/A | A comma-separated list of services that use this certificate. | N | none |
| **password** | String | The PKCS#12 file password. | Y | N/A |
| **name** | String | The certificate's symbolic name. | Y | N/A |
| **certificate** | N/A | The PKCS#12 content in one-line base64 format. Such input can be created, for example, by a base64 utility: base64 -w0 myCert.pfx | Y | N/A |

As a security precaution, use the **pki_show_security** command to view the certificate in plain text, and make sure that the certificate text under *Signature Algorithm* does not include the string *MD5*. This will help you avoid a "transcript collision" attack, that can force a hash-construction downgrade to MD5 and reduce expected security. For the vulnerability summary, see the National Vulnerability Database.

**Note:**

Quorum Witness communication might break in cases where replacing a certificate after HyperSwap relations have been created, if the default certificate has been overwritten while specifying the specific quorum service. This occurs when there are existing HyperSwap relations on the systems registered on the Quorum Witness and they are not deactived and before the certificate change. For more information, see the latest Quorum Witness user information within the FlashSystem A9000 or A9000R documentation set on IBM Documentation.

**Example:**

```
pki_set_pkcs12 name=myPki password=pkiPassword certificate=pkiCertificateBase64
```

**Output:**

```
Command completed successfully
```

## Access control

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

### Return codes

- **SERVICE_IS_USING_OTHER_CERTIFICATE**

  Service '*services*' is using another certificate.

  **Troubleshooting:** Edit the certificate previously used by the service.
- **CANNOT_VALIDATE_PKCS12_FILE**

  Failed validating PKCS#12 file.

  **Troubleshooting:** Make sure that the PKCS#12 file content is encoded to base64, and the password is correct.
- **DEFAULT_CERTIFICATE_ALREADY_EXIST**

  The default certificate already exists.

  **Troubleshooting:** Delete the default certificate or make it not default.
- **CERTIFICATE_NAME_ALREADY_EXIST**

  A certificate with the indicated name already exists.

  **Troubleshooting:** Enter a different name.
- **BAD_BASE64_DATA**

  Data cannot be decoded as base-64 data.
- **FAILED_GETTING_PRIVATE_KEY_FINGERPRINT**

  Failed to retrieve a private key fingerprint.
- **FAILED_ENCRYPTING_PRIVATE_KEY**

  Failed to encrypt a private key.
- **CERTIFICATE_CONTAINER_FULL**

  Cannot add any more certificates, the maximum number is already reached.

  **Troubleshooting:** Delete a certificate.

## Displaying the details of a signed certificate

Use the **pki_show_certificate** command to display the details of a signed certificate.

```
pki_show_certificate name=Name
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **name** | String | The certificate's symbolic name. | Y |

As a security precaution, use this command to view the certificate in plain text, and make sure that the certificate text under *Signature Algorithm* does not include the string *MD5*. This will help you avoid a "transcript collision" attack, that can force a hash-construction downgrade to MD5 and reduce expected security. For the vulnerability summary, see the National Vulnerability Database.

**Example:**

```
pki_show_certificate name=ibm
```

| Field ID | Field output | Default position |
|---|---|---|
| **certificate** | Certificate | 1 |

### Access control

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

### Return codes

- **FAILED_PARSING_CERTIFICATE**

  Failed parsing the certificate.
- **KEY_HAS_NO_CERTIFICATE**

  The key has no signed certificate defined.
- **CERTIFICATE_NAME_DOES_NOT_EXIST**

  A certificate with the indicated name was not found.

  **Troubleshooting:** Enter a different name.

## Updating a PKI certificate or services

Use the **pki_update** command to update a PKI certificate or services.

```
pki_update  name=Name [ services=<xcli [ ,cim ] [ ,ipsec ] ... | ALL | NONE> ]
 [ certificate=SignedCertificate ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **services** | N/A | Comma-separated list of services that need to use this certificate. | N | none |
| **name** | String | The certificate's symbolic name. | Y | N/A |
| **certificate** | N/A | If this parameter is defined, the certificate will be replaced. | N | none |

**Example:**

```
pki_update name=cert services=xcli,cim
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Allowed |

| User Category | Permission |
|---|---|
| Read-only users | Disallowed |
| Technicians | Disallowed |

## Return codes

- **SERVICE_IS_USING_OTHER_CERTIFICATE**

  Service '*services*' is using another certificate.

  **Troubleshooting:** Edit the certificate previously used by the service.
- **NO_PKI_UPDATE_PARAMETERS_SPECIFIED**

  No parameters were specified for the update.
- **CERTIFICATE_DOES_NOT_MATCH_PRIVATE_KEY**

  The certificate does not match the private key.

  **Troubleshooting:** Use another certificate.
- **CANNOT_SET_SERVICES_BEFORE_SETTING_CERTIFICATE**

  Cannot set services before setting the certificate.

  **Troubleshooting:** Set the certificate first.
- **DEFAULT_CERTIFICATE_ALREADY_EXIST**

  The default certificate already exists.

  **Troubleshooting:** Delete the default certificate or make it not default.
- **CERTIFICATE_KEY_WAS_NOT_FOUND**

  Failed to set the certificate.

  **Troubleshooting:** Make sure the certificate parameters are correct.
- **CERTIFICATE_NAME_DOES_NOT_EXIST**

  A certificate with the indicated name was not found.

  **Troubleshooting:** Enter a different name.
- **SSL_CERTIFICATE_CHAIN_EMPTY**

  No certificates were found in the input.
- **SSL_CERTIFICATE_HAS_EXPIRED**

  The SSL certificate has expired.
- **SSL_CERTIFICATE_INVALID_FORMAT**

  The SSL certificate format is invalid or corrupted.
- **SSL_CERTIFICATE_ISSUER_NOT_FOUND**

  The SSL certificate issuer was not found in the certificate chain.
- **SSL_CERTIFICATE_NOT_YET_VALID**

  The SSL certificate is not yet valid.
- **SSL_CERTIFICATE_VERIFICATION_FAILED**

  The SSL certificate chain verification failed.
- **SSL_CERTIFICATE_VERIFICATION_INTERNAL_ERROR**

  The SSL certificate verification has failed because of an internal system error.

# 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.is_cm_ok** | CM OK | N/A |
| **saved_info.port_down_reason** | Failure Reason | 7 |
| **saved_info.last_state_change** | Last State Change | N/A |
| **saved_info.last_cm_check** | Last CM Check | N/A |
| **pending_ia_cmd** | Component Operation | N/A |
| **currently_functioning** | Currently Functioning | N/A |
| **requires_service** | Requires Service | N/A |
| **service_reason** | Service Reason | N/A |

### Access control

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

© Copyright IBM Corp. 2016, 2022

# Listing data counters for the enabled InfiniBand HCA ports

Use the **ib_hca_counter_list** command to list data counters for InfiniBand HCA ports, enabled on modules and flash enclosures.

```
ib_hca_counter_list [ hca_port=ComponentId ]
```

## Parameters

| Name | Description | Mandatory | Default |
|---|---|---|---|
| **hca_port** | The InfiniBand HCA port to be listed. | N | All IB HCA ports |

**Example:**

```
ib_hca_counter_list
```

**Output:**

```
Port                          TX Data        RX Data        TX Pkt
----------------------------- -------------- -------------- ------------
1:IB_FlashSystem_Port:4:1     0              0              0
1:IB_FlashSystem_Port:4:3     0              0              0
1:IB_FlashSystem_Port:4:5     0              0              0
1:IB_FlashSystem_Port:4:7     0              0              0
1:IB_Module_Port:12:1         254584779527   252260909954   2055238854
1:IB_Module_Port:12:2         0              0              0
1:IB_Module_Port:13:1         252395242864   254798454598   2061534883
1:IB_Module_Port:13:2         0              0              0
1:IB_Module_Port:8:1          254003578209   254027205845   2055494787
1:IB_Module_Port:8:2          0              0              0

RX Pkt       XmtWait
------------ -----------
0            0
0            0
0            0
0            0
2058771428   103686442
0            0
2060782961   94235849
0            0
2059021166   103031319
0            0
```

| Field ID | Field output | Default position |
|---|---|---|
| **port** | Port | 1 |
| **XmtData** | TX Data | 2 |
| **RcvData** | RX Data | 3 |
| **XmtPkts** | TX Pkt | 4 |
| **RcvPkts** | RX Pkt | 5 |
| **XmtWait** | XmtWait | 6 |

## Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |

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

# Listing error counters for the enabled InfiniBand HCA ports.

Use the **ib_hca_error_list** command to list error counters for the enabled InfiniBand HCA ports.

```
ib_hca_error_list [ hca_port=ComponentId ]
```

## Parameters

| Name | Description | Mandatory | Default |
|---|---|---|---|
| **hca_port** | The InfiniBand HCA port to be listed. | N | All IB HCA ports |

### Example:

```
ib_hca_error_list
```

| Field ID | Field output | Default position |
|---|---|---|
| **port** | Port | 1 |
| **SymbolErrors** | SymErr | 2 |
| **LinkRecovers** | LinkRec | 3 |
| **LinkDowned** | LinkDown | 4 |
| **RcvErrors** | RcvErr | 5 |
| **RcvRemotePhysErrors** | RcvRPErr | 6 |
| **RcvSwRelayErrors** | RcvSRErr | 7 |
| **XmtDiscards** | XmtDisc | 8 |
| **XmtConstraintErrors** | XmtCErr | 9 |
| **RcvConstraintErrors** | RcvCErr | 10 |
| **LinkIntegrityErrors** | LinkIErr | 11 |
| **ExcBufOverrunErrors** | ExcBOEre | 12 |
| **VL15Dropped** | VL15Dr | 13 |

## Access control

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

# Listing the statuses of the enabled InfiniBand HCA ports

Use the **ib_hca_info_list** command to list the statuses of InfiniBand HCA ports, enabled on modules and flash enclosures.

```
ib_hca_info_list [ hca_port=ComponentId ]
```

## Parameters

| Name | Description | Mandatory | Default |
|---|---|---|---|
| **hca_port** | The InfiniBand HCA port to be listed. | N | All IB HCA ports |

**Example:**

```
ib_hca_info_list
```

**Output:**

```
Port                          IB Log State   IB Phys State
----------------------------- -------------- ---------------
1:IB_FlashSystem_Port:4:1   NOT SAMPLED    NOT SAMPLED
1:IB_FlashSystem_Port:4:3   NOT SAMPLED    NOT SAMPLED
1:IB_FlashSystem_Port:4:5   NOT SAMPLED    NOT SAMPLED
1:IB_FlashSystem_Port:4:7   NOT SAMPLED    NOT SAMPLED
1:IB_Module_Port:12:1       ACTIVE         LINK UP
1:IB_Module_Port:12:2       INIT           LINK UP
1:IB_Module_Port:13:1       ACTIVE         LINK UP
1:IB_Module_Port:13:2       INIT           LINK UP
1:IB_Module_Port:8:1        ACTIVE         LINK UP
1:IB_Module_Port:8:2        INIT           LINK UP

Link Speed     Link Width
-------------- -------------
NOT SAMPLED    NOT SAMPLED
NOT SAMPLED    NOT SAMPLED
NOT SAMPLED    NOT SAMPLED
NOT SAMPLED    NOT SAMPLED
14.0625 Gbps   X4
14.0625 Gbps   X4
14.0625 Gbps   X4
14.0625 Gbps   X4
14.0625 Gbps   X4
14.0625 Gbps   X4
```

| Field ID | Field output | Default position |
|---|---|---|
| **port** | Port | 1 |
| **log_state** | IB Log State | 2 |
| **phys_state** | IB Phys State | 3 |
| **link_speed** | Link Speed | 4 |
| **link_width** | Link Width | 5 |
| **link_width_sup** | Link Width Sup | N/A |
| **link_speed_sup** | Link Speed Sup | N/A |
| **link_speed_enabled** | Link Speed Ena | N/A |
| **link_width_enabled** | Link Width Ena | N/A |

### Access control

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

## Listing the configured InfiniBand switches

Use the **switch_list** command to list the configured InfiniBand switches.

```
switch_list [ switch=ComponentId ]
```

### Parameters

| Name | Description | Mandatory | Default |
|---|---|---|---|
| **switch** | IB switch to list. | N | All IB switches |

**Example:**

```
switch_list
```

**Output:**

```
Switch          Switch GUID        Status   MGMT OK   Ports OK   Power OK   BBU OK
--------------- ------------------ -------- --------- ---------- ---------- --------
1:IB_Switch:1   E41D2D03003C9900   OK       yes       yes        yes        yes
1:IB_Switch:2   E41D2D03003C9A80   OK       yes       yes        yes        yes

Cont.:

Fan OK   Temp OK   Volt OK   Boot Time             FW         Serial
-------- --------- --------- --------------------- ---------- --------------
yes      yes       yes       04/04/2016 10:51:18   3.5.0500   MT1523X09088
yes      yes       yes       27/03/2016 15:28:31   3.5.0500   MT1523X09091
```

| Field ID | Field output | Default position |
|---|---|---|
| **component_id** | Switch | 1 |
| **status** | Status | 3 |
| **sw_mgmt_status** | MGMT Status | N/A |
| **num_of_down_ports** | Down Ports | N/A |
| **mgmt_ok** | MGMT OK | 4 |
| **ports_ok** | Ports OK | 5 |
| **power_ok** | Power OK | 6 |
| **bbu_ok** | BBU OK | 7 |
| **fan_ok** | Fan OK | 8 |
| **temp_ok** | Temp OK | 9 |
| **volt_ok** | Volt OK | 10 |

| Field ID | Field output | Default position |
|---|---|---|
| **fw** | FW | 12 |
| **mgmt_serial_number** | Serial | 13 |
| **mgmt_part_number** | Part No | N/A |
| **mgmt_asic_rev** | ASIC Rev | N/A |
| **mgmt_hw_rev** | HW Rev | N/A |
| **cpld_tor** | CPLD Tor | N/A |
| **cpld_port1** | CPLD Port1 | N/A |
| **cpld_switch_brd** | CPLD Switch Brd | N/A |
| **chassis_serial_number** | Chassis Serial | N/A |
| **chassis_part_number** | Chassis Part No | N/A |
| **chassis_asic_rev** | Chassis ASIC Rev | N/A |
| **chassis_hw_rev** | Chassis HW Rev | N/A |
| **original_mgmt_serial_number** | Original Serial | N/A |
| **original_mgmt_part_number** | Original Part No | N/A |
| **original_mgmt_asic_rev** | Original ASIC Rev | N/A |
| **original_mgmt_hw_rev** | Original HW Rev | N/A |
| **original_chassis_serial_number** | Original Chassis Serial | N/A |
| **original_chassis_part_number** | Original Chassis Part No | N/A |
| **original_chassis_asic_rev** | Original Chassis ASIC Rev | N/A |
| **original_chassis_hw_rev** | Original Chassis HW Rev | N/A |
| **currently_functioning** | Currently Functioning | N/A |
| **mgmt_guid** | Managememt GUID | N/A |
| **requires_service** | Requires Service | N/A |
| **service_reason** | Service Reason | N/A |
| **fabric_switch_info.miss_count** | Miss Count | N/A |
| **fabric_switch_info.num_of_ports** | Ports | N/A |
| **fabric_switch_info.image_guid** | FW GUID | N/A |
| **fabric_switch_info.node_guid** | Switch GUID | 2 |
| **fabric_switch_info.dev_id** | Device ID | N/A |
| **fabric_switch_info.dev_rev** | Device Revision | N/A |
| **fabric_switch_info.vendor_id** | Vendor ID | N/A |
| **fabric_switch_info.name** | Name | N/A |
| **fabric_switch_info.mlx_dev_id** | Ext Device ID | N/A |
| **fabric_switch_info.mlx_hw_rev** | Ext Device Revision | N/A |
| **fabric_switch_info.boot_time** | Boot Time | 11 |
| **fabric_switch_info.uptime_seconds** | Uptime | N/A |
| **fabric_switch_info.fw_build_id** | FW BUILD ID | N/A |
| **fabric_switch_info.fw_rev** | FW Ver | N/A |
| **fabric_switch_info.fw_build_date** | FW Build Date | N/A |

| Field ID | Field output | Default position |
|---|---|---|
| **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 |

| Field ID | Field output | Default position |
|---|---|---|
| **router** | Router | 4 |
| **real_ip** | Real IP | 5 |
| **status** | Status | 6 |

### Access control

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

## Listing the configured InfiniBand switch firmware versions

Use the **switch_fw_list** command to list the configured InfiniBand switch firmware versions.

```
switch_fw_list [ switch=ComponentId ]
```

### Parameters

| Name | Description | Mandatory | Default |
|---|---|---|---|
| **switch** | IB switch to list. | N | All IB switches |

**Example:**

```
switch_fw_list
```

**Output:**

```
Switch          Type              Version
--------------- ----------------- ----------
1:IB_Switch:1   ASIC              9.3.7170
1:IB_Switch:1   BBU1              703
1:IB_Switch:1   BBU2              703
1:IB_Switch:1   BIOS              4.6.5
1:IB_Switch:1   CPLD_PORT1        4
1:IB_Switch:1   CPLD_SWITCH_BRD   7
1:IB_Switch:1   CPLD_TOR          9
1:IB_Switch:1   MGMT              3.5.0500
1:IB_Switch:1   PSU1              404
1:IB_Switch:1   PSU2              404
1:IB_Switch:2   ASIC              9.3.7170
1:IB_Switch:2   BBU1              703
1:IB_Switch:2   BBU2              703
1:IB_Switch:2   BIOS              4.6.5
1:IB_Switch:2   CPLD_PORT1        4
1:IB_Switch:2   CPLD_SWITCH_BRD   7
1:IB_Switch:2   CPLD_TOR          9
1:IB_Switch:2   MGMT              3.5.0500
1:IB_Switch:2   PSU1              404
1:IB_Switch:2   PSU2              404
```

| Field ID | Field output | Default position |
|---|---|---|
| **switch_id** | Switch | 1 |
| **type** | Type | 2 |
| **version** | Version | 3 |
| **original_version** | Original Version | N/A |

### Access control

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

## Listing the configured InfiniBand switch power values

Use the **switch_power_list** command to list the configured InfiniBand switch power values for PSUs and BBUs.

```
switch_power_list [ switch=ComponentId ]
```

### Parameters

| Name | Description | Mandatory | Default |
|---|---|---|---|
| **switch** | IB switch to list. | N | All IB switches |

**Example:**

```
switch_power_list
```

**Output:**

```
Switch          Type   Power W   Voltage V   Current A   Capacity W   Feed   Status
--------------- ------ --------- ----------- ----------- ------------ ------ --------
1:IB_Switch:1   BBU1   N/A       N/A         N/A         330.00       NA     OK
1:IB_Switch:1   BBU2   N/A       N/A         N/A         330.00       NA     OK
1:IB_Switch:1   PS1    46.00     12.11       2.56        400.00       AC     OK
1:IB_Switch:1   PS2    44.00     12.19       2.75        400.00       AC     OK
1:IB_Switch:2   BBU1   N/A       N/A         N/A         330.00       NA     OK
1:IB_Switch:2   BBU2   N/A       N/A         N/A         330.00       NA     OK
1:IB_Switch:2   PS1    47.00     12.05       2.56        400.00       AC     OK
1:IB_Switch:2   PS2    44.00     12.14       2.81        400.00       AC     OK
```

| Field ID | Field output | Default position |
|---|---|---|
| **switch_id** | Switch | 1 |
| **type** | Type | 2 |
| **power** | Power W | 3 |
| **voltage** | Voltage V | 4 |
| **current** | Current A | 5 |

| Field ID | Field output | Default position |
|---|---|---|
| **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 |

**Note:**

The temperature values are indicated in Celsius.

**Example:**

```
switch_temp_list
```

**Output:**

```
Switch          Type         Actual   Alert    Critical   Status
--------------- ------------ -------- -------- ---------- --------
1:IB_Switch:1   BBU1         19.60    60.00    60.00      OK
1:IB_Switch:1   BBU2         19.50    60.00    60.00      OK
1:IB_Switch:1   MGMT_AMB     20.50    120.00   120.00     OK
1:IB_Switch:1   MGMT_CPU     25.00    120.00   120.00     OK
1:IB_Switch:1   MGMT_CPU1    22.00    120.00   120.00     OK
1:IB_Switch:1   MGMT_CPU2    25.00    120.00   120.00     OK
1:IB_Switch:1   MGMT_PORTS   22.00    120.00   120.00     OK
1:IB_Switch:1   MGMT_SX      28.00    105.00   110.00     OK
1:IB_Switch:1   PS1          24.00    120.00   120.00     OK
1:IB_Switch:1   PS2          24.00    120.00   120.00     OK
1:IB_Switch:2   BBU1         20.00    60.00    60.00      OK
1:IB_Switch:2   BBU2         19.89    60.00    60.00      OK
1:IB_Switch:2   MGMT_AMB     21.00    120.00   120.00     OK
1:IB_Switch:2   MGMT_CPU     26.00    120.00   120.00     OK
1:IB_Switch:2   MGMT_CPU1    25.00    120.00   120.00     OK
1:IB_Switch:2   MGMT_CPU2    19.00    120.00   120.00     OK
1:IB_Switch:2   MGMT_PORTS   22.50    120.00   120.00     OK
1:IB_Switch:2   MGMT_SX      28.00    105.00   110.00     OK
1:IB_Switch:2   PS1          24.00    120.00   120.00     OK
1:IB_Switch:2   PS2          25.00    120.00   120.00     OK
```

| Field ID | Field output | Default position |
|---|---|---|
| **switch_id** | Switch | 1 |
| **type** | Type | 2 |
| **actual** | Actual | 3 |
| **alert** | Alert | 4 |
| **critical** | Critical | 5 |
| **status** | Status | 6 |

## Access control

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

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

# Listing the configured InfiniBand switch fan parts

Use the **switch_fan_part_list** command to list the configured InfiniBand switch fan parts.

```
switch_fan_part_list [ switch=ComponentId ]
```

## Parameters

| Name | Description | Mandatory | Default |
|---|---|---|---|
| **switch** | IB switch to list. | N | All IB switches |

**Example:**

```
switch_fan_part_list
```

**Output:**

```
Switch          Type      Part No          Serial No      HW Rev   Speed      Status
--------------- --------- ---------------- -------------- -------- ---------- --------
1:IB_Switch:1   FAN1-F1   MTEF-FANF-B      MT1523X09083   A2       10608.00   OK
1:IB_Switch:1   FAN1-F2   MTEF-FANF-B      MT1523X09083   A2       8998.00    OK
1:IB_Switch:1   FAN2-F1   MTEF-FANF-B      MT1523X09075   A2       10526.00   OK
1:IB_Switch:1   FAN2-F2   MTEF-FANF-B      MT1523X09075   A2       8939.00    OK
1:IB_Switch:1   FAN3-F1   MTEF-FANF-B      MT1523X09082   A2       10608.00   OK
1:IB_Switch:1   FAN3-F2   MTEF-FANF-B      MT1523X09082   A2       8998.00    OK
1:IB_Switch:1   FAN4-F1   MTEF-FANF-B      MT1523X09077   A2       10691.00   OK
1:IB_Switch:1   FAN4-F2   MTEF-FANF-B      MT1523X09077   A2       9242.00    OK
1:IB_Switch:1   PSU1      MTEF-PSF-AC-B    MT1523X09040   A5       14464.00   OK
1:IB_Switch:1   PSU2      MTEF-PSF-AC-B    MT1523X09044   A5       14592.00   OK
1:IB_Switch:2   FAN1-F1   MTEF-FANF-B      MT1523X09065   A2       10526.00   OK
1:IB_Switch:2   FAN1-F2   MTEF-FANF-B      MT1523X09065   A2       9118.00    OK
1:IB_Switch:2   FAN2-F1   MTEF-FANF-B      MT1523X09072   A2       10445.00   OK
1:IB_Switch:2   FAN2-F2   MTEF-FANF-B      MT1523X09072   A2       9118.00    OK
1:IB_Switch:2   FAN3-F1   MTEF-FANF-B      MT1523X09062   A2       10445.00   OK
1:IB_Switch:2   FAN3-F2   MTEF-FANF-B      MT1523X09062   A2       8998.00    OK
1:IB_Switch:2   FAN4-F1   MTEF-FANF-B      MT1523X09073   A2       10526.00   OK
1:IB_Switch:2   FAN4-F2   MTEF-FANF-B      MT1523X09073   A2       8998.00    OK
1:IB_Switch:2   PSU1      MTEF-PSF-AC-B    MT1523X09035   A5       14464.00   OK
1:IB_Switch:2   PSU2      MTEF-PSF-AC-B    MT1523X09033   A5       14464.00   OK
```

| Field ID | Field output | Default position |
|---|---|---|
| **switch_id** | Switch | 1 |
| **type** | Type | 2 |
| **part_no** | Part No | 3 |
| **serial_no** | Serial No | 4 |
| **hw_rev** | HW Rev | 5 |
| **speed** | Speed | 6 |
| **status** | Status | 7 |

### Access control

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

## Listing the configured InfiniBand switch PSUs

Use the **switch_psu_list** command to list the configured InfiniBand switch PSUs.

```
switch_psu_list [ switch_psu=ComponentId ] [ switch=ComponentId ]
```

### Parameters

| Name | Description | Mandatory | Default |
|---|---|---|---|
| **switch_psu** | IB switch PSU to list. | N | All IB switch PSUs |
| **switch** | IB switch whose BBUs are to be listed. | N | All IB switches |

**Example:**

```
switch_psu_list
```

**Output:**

```
Component ID          Status   Sensor Status   Power W   Voltage V
--------------------- -------- --------------- --------- -----------
1:IB_Switch_PSU:1:1   OK       OK              46.00     12.11
1:IB_Switch_PSU:1:2   OK       OK              44.00     12.19
1:IB_Switch_PSU:2:1   OK       OK              47.00     12.05
1:IB_Switch_PSU:2:2   OK       OK              44.00     12.14

Cont.:

Current A   Capacity W
----------- ------------
2.94        400.00
2.50        400.00
2.69        400.00
3.25        400.00
```

| Field ID | Field output | Default position |
|---|---|---|
| **component_id** | Component ID | 1 |
| **status** | Status | 2 |
| **sensor_status** | Sensor Status | 3 |
| **power** | Power W | 4 |
| **voltage** | Voltage V | 5 |
| **current** | Current A | 6 |
| **capacity** | Capacity W | 7 |

| Field ID | Field output | Default position |
|---|---|---|
| **temperature** | Temperature | N/A |
| **fan_speed** | Fan Speed | N/A |
| **fan_sensor_status** | Fan Status | N/A |
| **serial_number** | Serial No | N/A |
| **original_serial_number** | Original Serial No | N/A |
| **part_number** | Part No | N/A |
| **original_part_number** | Original Part No | N/A |
| **hw_rev** | HW Rev | N/A |
| **requires_service** | Requires Service | N/A |
| **service_reason** | Service Reason | N/A |
| **currently_functioning** | Currently Functioning | N/A |
| **switch_id** | Switch ID | N/A |

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Allowed |
| Security administrator | Disallowed |
| Read-only users | Allowed |
| Technicians | Allowed |

## Listing the configured InfiniBand switch BBUs

Use the **switch_bbu_list** command to list the configured InfiniBand switch BBUs.

```
switch_bbu_list [ switch_bbu=ComponentId ] [ switch=ComponentId ]
```

### Parameters

| Name | Description | Mandatory | Default |
|---|---|---|---|
| **switch_bbu** | IB switch BBU to list. | N | All IB switch BBUs |
| **switch** | IB switch whose BBUs are to be listed. | N | All IB switches |

**Example:**

```
switch_bbu_list
```

**Output:**

```
Component ID          Status   Sensor Status   Remaining capacity   Full charged capacity
--------------------- -------- --------------- -------------------- ----------------------
1:IB_Switch_BBU:1:1   OK       OK              71520 mWh            80660 mWh
1:IB_Switch_BBU:1:2   OK       OK              46060 mWh            55520 mWh
1:IB_Switch_BBU:2:1   OK       OK              80800 mWh            80800 mWh
1:IB_Switch_BBU:2:2   OK       OK              77840 mWh            80120 mWh

Cont.:

Percent Charged   Charger Status   Calibration Status   Last Recondition Date
----------------- ---------------- -------------------- ----------------------
100%              Fully charged    N/A                  N/A
100%              Fully charged    N/A                  N/A
100%              Fully charged    N/A                  N/A
100%              Fully charged    N/A                  N/A
```

| Field ID | Field output | Default position |
|---|---|---|
| **component_id** | Component ID | 1 |
| **status** | Status | 2 |
| **sensor_status** | Sensor Status | 3 |
| **relative_capacity** | Remaining capacity | 4 |
| **absolute_capacity** | Full charged capacity | 5 |
| **relative_capacity_percent** | Percent Charged | 6 |
| **charging_state** | Charger Status | 7 |
| **remaining_charging_time** | Remaining Charging Time | N/A |
| **calibration_status** | Calibration Status | 8 |
| **last_calibration_date** | Last Recondition Date | 9 |
| **fw_version** | FW | N/A |
| **serial_number** | Serial No | N/A |
| **original_serial_number** | Original Serial No | N/A |
| **part_number** | Part No | N/A |
| **original_part_number** | Original Part No | N/A |
| **voltage** | Voltage | N/A |
| **exp_voltage** | Expected Voltage | N/A |
| **temperature** | Temperature | N/A |
| **manufacture_date** | Manufacture Date | N/A |
| **designed_capacity** | Designed Capacity | N/A |
| **absolute_charge** | Absolute Charge | N/A |
| **test_status** | Test Status | N/A |
| **last_test_date** | Last Test Date | N/A |
| **can_charge** | Charge | N/A |
| **can_discharge** | Discharge | N/A |
| **charge_rate** | Charge Rate | N/A |
| **max_error** | Max Error | N/A |
| **hw_rev** | HW Rev | N/A |
| **requires_service** | Requires Service | N/A |
| **service_reason** | Service Reason | N/A |
| **currently_functioning** | Currently Functioning | N/A |

| Field ID | Field output | Default position |
|---|---|---|
| **switch_id** | Switch ID | N/A |
| **active_alarms** | Active Alarms | N/A |

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Allowed |
| Security administrator | Disallowed |
| Read-only users | Allowed |
| Technicians | Allowed |

## Listing the configured InfiniBand switch fans

Use the **switch_fan_list** command to list the configured InfiniBand switch fans.

```
switch_fan_list [ switch_fan=ComponentId ] [ switch=ComponentId ]
```

### Parameters

| Name | Description | Mandatory | Default |
|---|---|---|---|
| **switch_fan** | IB switch fan to list. | N | All IB switch fans |
| **switch** | IB switch to list the BBUs of. | N | All IB switches |

### Example:

```
switch_fan_list
```

### Output:

```
Component ID          Status   Speed      Sensor Status   Peer Speed   Peer Sensor Status
--------------------- -------- ---------- --------------- ------------ --------------------
1:IB_Switch_Fan:1:1   OK       10608.00   OK              8998.00      OK
1:IB_Switch_Fan:1:2   OK       10526.00   OK              8939.00      OK
1:IB_Switch_Fan:1:3   OK       10608.00   OK              8998.00      OK
1:IB_Switch_Fan:1:4   OK       10691.00   OK              9242.00      OK
1:IB_Switch_Fan:2:1   OK       10526.00   OK              9118.00      OK
1:IB_Switch_Fan:2:2   OK       10445.00   OK              9118.00      OK
1:IB_Switch_Fan:2:3   OK       10445.00   OK              8998.00      OK
1:IB_Switch_Fan:2:4   OK       10526.00   OK              8998.00      OK
```

| Field ID | Field output | Default position |
|---|---|---|
| **component_id** | Component ID | 1 |
| **status** | Status | 2 |
| **speed** | Speed | 3 |
| **sensor_status** | Sensor Status | 4 |
| **peer_speed** | Peer Speed | 5 |
| **peer_sensor_status** | Peer Sensor Status | 6 |

| Field ID | Field output | Default position |
|---|---|---|
| **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 |

# 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 completed successfully.
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Disallowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

### Return codes

- **USER_GROUP_NAME_DOES_NOT_EXIST**

  The user group name does not exist.
- **CLUSTER_BAD_NAME**

  The cluster name does not exist.
- **HOST_BAD_NAME**

  The host name does not exist.

© Copyright IBM Corp. 2016, 2022

- **HOST_BELONGS_TO_CLUSTER**

  This host already belongs to a cluster.

# Deleting an access control definition

Use the **access_delete** command to delete an access control definition.

```
access_delete user_group=UserGroup < host=HostName | cluster=ClusterName >
```

## Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **user_group** | Object name | The user group specified in the access control definition that should be deleted. | Y |
| **host** | Object name | The host specified in the access control definition that should be deleted. | N |
| **cluster** | Object name | The cluster specified in the access control definition that should be deleted. | N |

This command deletes an association between the user group and host or cluster. The operation fails if no such access definition exists. When a host is removed from a cluster, the host's associations become the cluster's associations. This allows a continued mapping of operations, so that all scripts continue to work.

**Example:**

```
access_delete user_group=usergroup1
```

**Output:**

```
Command completed successfully.
```

## Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Disallowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

## Return codes

- **USER_GROUP_NAME_DOES_NOT_EXIST**

  The user group name does not exist.
- **USER_GROUP_DOES_NOT_HAVE_ACCESS_TO_CLUSTER**

  User Group does not have access to cluster.
- **CLUSTER_BAD_NAME**

  The cluster name does not exist.

- **HOST_BAD_NAME**

  The host name does not exist.
- **HOST_BELONGS_TO_CLUSTER**

  This host already belongs to a cluster.
- **USER_GROUP_DOES_NOT_HAVE_ACCESS_TO_HOST**

  User Group does not have access to host.

## Listing access control definitions

Use the **access_list** command to list access control definitions.

```
access_list [ user_group=UserGroup ] [ host=HostName | cluster=ClusterName ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **user_group** | Object name | Filters the access control listing to display only this user group. | N | All user groups. |
| **host** | Object name | Filters the access control listing to display only this host. | N | All hosts. |
| **cluster** | Object name | Filters the access control listing to display only this cluster. | N | All clusters. |

The list can be displayed for all access control definitions, or it can be filtered for a specific user group, host/cluster, or both.

| Field ID | Field output | Default position |
|---|---|---|
| **type** | Type | 1 |
| **name** | Name | 2 |
| **user_group** | User Group | 3 |

**Example:**

```
access_list host=buffyvam
```

**Output:**

```
Type   Name       User Group
host   buffyvam   testing
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Allowed |
| Security administrator | Disallowed |
| Read-only users | Allowed |

| User Category | Permission |
|---|---|
| Technicians | Disallowed |

### Return codes

- **HOST_BAD_NAME**

  The host name does not exist.
- **CLUSTER_BAD_NAME**

  The cluster name does not exist.
- **USER_GROUP_NAME_DOES_NOT_EXIST**

  The user group name does not exist.

## Adding an LDAP server definition

Use the **ldap_add_server** command to add an LDAP server definition.

```
ldap_add_server  fqdn=Fqdn  [ address=Address ] base_dn=LdapDn [ certificate=PemCertificate ]
[ port=PortNum ]
 [ secure_port=PortNum ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **fqdn** | N/A | FQDN of the LDAP server. | Y | N/A |
| **address** | N/A | IP address of the LDAP server. Only required when DNS is not available for FQDN to IP address resolution. | N | none |
| **base_dn** | N/A | Base_DN of the LDAP server. Serves as the starting reference point for searches. | Y | N/A |
| **certificate** | N/A | The public certificate or certificate chain of the LDAP server to be added (see below for details). | N | no certificate |
| **port** | Integer | The port number. | N | 389 |
| **secure_port** | Integer | The secure port number. | N | 636 |

#### The certificate parameter

The value of the certificate parameter is the content of a PEM file with asterisks instead of newlines. Chained certificates are supported. The total maximal length of a PEM file holding chained certificates (leaf first, root last) is 15360 characters (including the asterisk characters). In Windows, you can drag-and-drop a PEM file from the Windows Explorer to the appropriate location in the XCLI session window; the content will be added automatically.

**Example:**

```
ldap_add_server fqdn=ldap.example.com address=192.0.2.1
```

**Output:**

```
Command completed successfully.
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Disallowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

### Return codes

- **MAX_LDAP_SERVERS_REACHED**

  The number of defined LDAP servers already reached the maximum.
- **ADDRESS_CURRENTLY_ASSOCIATED_WITH_ANOTHER_LDAP_SERVER**

  The specified IP address is currently associated with another LDAP server.
- **LDAP_SERVER_EXISTS**

  LDAP server with the specified FQDN already exists.
- **SSL_CERTIFICATE_CHAIN_EMPTY**

  No certificates were found in the input.
- **SSL_CERTIFICATE_HAS_EXPIRED**

  The SSL certificate has expired.
- **SSL_CERTIFICATE_INVALID_FORMAT**

  The SSL certificate format is invalid or corrupted.
- **SSL_CERTIFICATE_ISSUER_NOT_FOUND**

  The SSL certificate issuer was not found in the certificate chain.
- **SSL_CERTIFICATE_NOT_YET_VALID**

  The SSL certificate is not yet valid.
- **SSL_CERTIFICATE_VERIFICATION_FAILED**

  The SSL certificate chain verification failed.
- **SSL_CERTIFICATE_VERIFICATION_INTERNAL_ERROR**

  The SSL certificate verification has failed because of an internal system error.

## Testing an LDAP configuration

Use the **ldap_test** command to authenticate the specified user against an LDAP server, based on the existing configuration.

```
ldap_test [ fqdn=Fqdn ] user=UserName    password=Password
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **fqdn** | N/A | FQDN of an LDAP server. | N | All servers |
| **user** | String | The username of the tested user. | Y | N/A |
| **password** | String | The password of the tested user. | Y | N/A |

**Example:**

```
ldap_test user=user1 password=pass1
```

**Output:**

```
Command completed successfully
```

## Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Disallowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

## Return codes

- **NO_LDAP_SERVERS_ARE_DEFINED**

  No LDAP servers are defined in the system
- **LDAP_SERVER_NOT_DEFINED**

  LDAP server *Server FQDN* is not defined in the system.
- **LDAP_IS_NOT_FULLY_CONFIGURED**

  LDAP is not fully configured.

  **Troubleshooting:** Check LDAP settings.
- **NO_LDAP_SERVERS_WITH_CERTIFICATE_ARE_DEFINED**

  No LDAP servers with an LDAP certificate are defined in the system.
- **SSL_CERTIFICATE_HAS_EXPIRED_FOR_SERVER**

  SSL certificate of LDAP server '*Server FQDN*' expired on *Expiration Date*.
- **USER_IS_PREDEFINED_IN_THE_SYSTEM**

  The user is predefined in the system.
- **LOGIN_FAILURE_USER_CANNOT_BE_UNIQUELY_AUTHENTICATED_BY_LDAP_SERVER**

  User *User Name* was not uniquely authenticated by LDAP server '*Server FQDN*'.
- **LOGIN_FAILURE_LDAP_SERVER_UNREACHABLE**

  No LDAP server can be reached.
- **LOGIN_FAILURE_XIV_USER_NOT_AUTHENTICATED_BY_LDAP_SERVER**

XIV User '*XIV User*' was not authenticated by LDAP server '*Server FQDN*'.

- **LOGIN_FAILURE_LDAP_SERVER_UNREACHABLE_OR_USER_NOT_FOUND**

  User *User Name* was not found in LDAP servers '*Servers FQDN*'.
- **LOGIN_FAILURE_INVALID_BASE_DN**

  The base dn of server '*Server FQDN*' is invalid.
- **LOGIN_FAILURE_USER_NOT_AUTHENTICATED_BY_LDAP_SERVER**

  User *User Name* was not authenticated by LDAP server '*Server FQDN*'.
- **LOGIN_FAILURE_USER_HAS_NO_RECOGNIZED_ROLE**

  User *User Name* has no recognized LDAP role.
- **LOGIN_FAILURE_USER_HAS_MORE_THAN_ONE_RECOGNIZED_ROLE**

  User *User Name* has more than one recognized LDAP role.
- **LOGIN_FAILURE_USER_MISSING_ID_ATTRIBUTE**

  User *User Name* is missing the LDAP ID attribute '*Attribute*'.
- **LOGIN_FAILURE_USER_MISSING_GROUP_ATTRIBUTE**

  User *User Name* is missing the group attribute '*Attribute*'.
- **LOGIN_FAILURE_USER_NOT_FOUND_IN_LDAP_SERVERS**

  User *User Name* was not found in LDAP servers.
- **LDAP_ROLE_UNRECOGNIZED**

  The LDAP role for the user is not recognized in the system.
- **LDAP_SERVER_NOT_FOUND**

  No LDAP server with the specified FQDN is defined in the system.
- **LDAP_AUTHENTICATION_IS_NOT_ACTIVE**

  LDAP authentication is not active.

## Listing LDAP configuration parameters

Use the **ldap_config_get** command to display system parameters that control user authentication against a specified LDAP server.

```
ldap_config_get
```

A successful execution of this command depends on connecting to a valid LDAP server.

The output of the command does not list LDAP servers. For the list of LDAP servers, use the **ldap_list_servers** command.

The **xiv_password** parameter is not listed.

**Example:**

```
ldap_config_get
```

**Output:**

```
Name                     Value
current_server
base_dn
version                  3
xiv_group_attrib
storage_admin_role
read_only_role
session_cache_period     20
bind_time_limit          20
user_id_attrib           objectSiD
first_expiration_event   30
second_expiration_event  14
third_expiration_event   7
use_ssl                  no
xiv_user
```

| Field ID | Field output | Default position |
|---|---|---|
| **name** | Name | 1 |
| **value** | Value | 2 |

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Disallowed |
| Read-only users | Allowed |
| Technicians | Disallowed |

## Configuring LDAP in the system

Use the **ldap_config_set** command to configure general system parameters that control user authentication against LDAP servers.

```
ldap_config_set [ user_name_attrib=LdapAttrib ] [
       xiv_group_attrib=LdapAttrib ]
 [ storage_admin_role=LdapRole ] [
       read_only_role=LdapRole ]
 [ security_admin_role=LdapRole ] [
       storage_integration_admin_role=LdapRole ]
 [
       xiv_host_profiler_role=LdapRole ] [ use_ssl=<yes|no> ]
 [
       user_id_attrib=LdapAttrib ] [ session_cache_period=Minutes ]
 [
       bind_time_limit=Seconds ] [ first_expiration_event=Days ]
 [
       second_expiration_event=Days ] [ third_expiration_event=Days ]
 [
       version=LdapVersion ] [ xiv_user=LdapAttrib ] [
       xiv_password=LdapAttrib ]
 [ server_type=<SUN DIRECTORY|MICROSOFT
       ACTIVE DIRECTORY|OPEN LDAP> ]
 [ group_search_depth=Depth ] [
       group_search_max_queries=Number ]
 [
       group_search_stop_when_found=<yes|no> ]
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **user_name_attrib** | String | User name attribute for queries. If not specified, it is set to `uid` for SUN Directory servers and `userPrincipalName` for Microsoft Active Directory servers. | N | According to server type |
| **xiv_group_attrib** | String | LDAP attribute designated to hold system-mapped roles. | N | none |
| **storage_admin_role** | String | LDAP value mapped to the Storage Administrator role. Multiple (up to 8) values are supported and must be separated with a semicolon (;). Multiple roles are not available for SUN Directory LDAP Servers. Up to 499 characters are allowed. | N | none |
| **read_only_role** | String | LDAP value mapped to the Read Only role. Multiple (up to 8) values are supported and must be separated with a semicolon (;). Up to 499 characters are allowed. | N | none |
| **security_admin_role** | String | LDAP value mapped to the Security Administrator role. Multiple (up to 8) values are supported and must be separated with a semicolon (;). Up to 499 characters are allowed. | N | none |
| **storage_integration_admin_role** | String | LDAP value mapped to the Storage Integration Administrator role. Multiple (up to 8) values are supported and must be separated with a semicolon (;). Up to 499 characters are allowed. | N | none |
| **xiv_host_profiler_role** | String | LDAP value mapped to the xiv host profiler role. Multiple (up to 8) values are supported and must be separated with a semicolon (;). Up to 499 characters are allowed. | N | none |
| **use_ssl** | Boolean | Indicates whether secure LDAP is mandatory. | N | no |
| **user_id_attrib** | String | The LDAP attribute set to identify the user (in addition to user DN) when recording user operations in the event log. | N | objectSiD |

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **session_cache_period** | Positive integer | Duration of keeping user credentials before attempting to re-login the user. | N | 20 |
| **bind_time_limit** | Positive integer | The duration after which the next LDAP server on the LDAP server list will be called. | N | 0. If set to the default, the LDAP server is contacted for every command. Performance issues depend on its availability. |
| **first_expiration_event** | Positive integer | The number of days before the expiration of certificate, when the first alert is issued (severity: warning). | N | 30/14/7 (third is smallest) |
| **second_expiration_event** | Positive integer | The number of days before the expiration of certificate, when the second alert is issued (severity: warning). | N | 30/14/7 (third is smallest) |
| **third_expiration_event** | Positive integer | The number of days before the expiration of certificate, when the third alert is issued (severity: warning). | N | 30/14/7 (third is smallest) |
| **version** | Positive integer | Version of LDAP used (only version 3 is supported). | N | 3 |
| **xiv_user** | String | The user for LDAP queries. | N | none |
| **xiv_password** | String | The password of user for LDAP queries. | N | none |
| **server_type** | Enumeration | Type of the LDAP server. | N | none |
| **group_search_depth** | Positive integer | The depth of group hierarchy to search in. | N | 0 |
| **group_search_max_queries** | Positive integer | Maximum number of group queries to perform per server. | N | 39 |
| **group_search_stop_when_found** | Boolean | Stop the group search when a group match is found. | N | yes |

LDAP access permissions are not enforced for predefined users. These predefined users are authenticated by the IBM storage system and not by LDAP even if LDAP authentication is enabled.

Predefined user names are:

- admin
- technician
- xiv_development
- xiv_maintenance

When an LDAP user, whose user name is identical with a predefined name, attempts to log into the system with LDAP authentication enabled, access will normally be denied, because:

- the user is not authenticated against LDAP, but rather against the storage system
- the user's (LDAP) password most likely does not match the storage system password.

However, if the user attempts to log into the system using the password of the corresponding predefined user, he or she will be granted the rights of the corresponding predefined user regardless of LDAP settings (for example, the user's association with the Application Administrator role), because LDAP authentication for predefined users is not required.

When an LDAP user, with a user name identical to any of the predefined names, tries to log into the system (when LDAP authentication is enabled), that user will normally be denied access, since he or she is not authenticated against LDAP (but rather against the storage system), and the user's (LDAP) password most likely does not match the storage system password. However, logging into the system using the password of the corresponding predefined user, the user will be granted the rights of the corresponding predefined user, regardless of LDAP settings (for example, association with an Application Administrator role), because LDAP is not consulted for predefined users.

**Example:**

```
ldap_config_set
            storage_admin_role="CN=EXMPLEAdmins,OU=EXMPLELab,DC=CROSSQA,DC=COM"
            re
ad_only_role="CN=EXMPLEOpers,OU=EXMPLELab,DC=CROSSQA,DC=COM"
            user_id_attrib=sAMA
ccountName user_name_attrib=sAMAccountName
            xiv_group_attrib=memberOf

 xiv_user="CN=allmighty,CN=Users,DC=CROSSQA,DC=COM"
            xiv_password=the_password

          security_admin_role="CN=EXMPLESecAdmins,OU=EXMPLELab,DC=CROSSQA,DC=COM"
```

**Output:**

```
Command completed successfully.
```

## Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Disallowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

## Return codes

- **LDAP_IS_NOT_FULLY_CONFIGURED**

  LDAP is not fully configured.

  **Troubleshooting:** Check LDAP settings.
- **LDAP_CONFIG_CHANGE_IS_ILLEGAL_WHEN_AUTHENTICATION_IS_ACTIVE**

  This LDAP configuration change is invalid when LDAP-based authentication is active.

  **Troubleshooting:** Disable the LDAP-based authentication and then change the LDAP configuration.
- **LDAP_ROLE_ALREADY_USED**

  This LDAP role is already in use in the LDAP configuration or in a user group.
- **NO_LDAP_SERVERS_WITH_CERTIFICATE_ARE_DEFINED**

  No LDAP servers with an LDAP certificate are defined in the system.

- **INVALID_EXPIRATION_EVENT_DATES**

  The dates for expiration events must be in ascending order.
- **LDAP_READ_ONLY_ROLE_HAS_WRONG_NUMBER_OF_PARTS**

  Too many parts in the LDAP read-only role.

  **Troubleshooting:** Role parts are divided by the ';' symbol. The number of parts should be between 0 and 8.
- **LDAP_ROLE_HAS_DUPLICATED_PARTS**

  The LDAP role contains duplicated parts.
- **LDAP_STORAGE_ADMIN_ROLE_HAS_WRONG_NUMBER_OF_PARTS**

  Too many parts in the LDAP storage administrator role.

  **Troubleshooting:** Role parts are divided by the ';' symbol. The number of parts should be between 0 and 8.
- **LDAP_SECURITY_ADMIN_ROLE_HAS_WRONG_NUMBER_OF_PARTS**

  Too many parts in the LDAP security administrator role.

  **Troubleshooting:** Role parts are divided by the ';' symbol. The number of parts should be between 0 and 8.
- **LDAP_STORAGE_INTEGRATION_ADMIN_ROLE_HAS_WRONG_NUMBER_OF_PARTS**

  Too many parts in the LDAP storage integration administrator role.

  **Troubleshooting:** Role parts are divided by the ';' symbol. The number of parts should be between 0 and 8.
- **LDAP_XIV_HOST_PROFILER_ROLE_HAS_WRONG_NUMBER_OF_PARTS**

  Too many parts in the LDAP XIV host profiler role.

  **Troubleshooting:** Role parts are divided by the ';' symbol. The number of parts should be between 0 and 8.

# Listing LDAP servers defined in the system

Use the **ldap_list_servers** command to list LDAP servers defined in the system.

```
ldap_list_servers  [ fqdn=Fqdn ]
```

## Parameters

| Name | Description | Mandatory | Default |
|---|---|---|---|
| **fqdn** | FQDN of a specific server to list. | N | All servers. |

This command lists the LDAP servers defined in the system along with their type description and the indication whether they are mandatory.

**Example:**

```
ldap_list_servers fqdn
```

**Output:**

```
<code value="SUCCESS"/>
        <empty_table_message value="No LDAP servers are defined in the system"/>
        <last_change_index value="367896"/>
        <status value="0"/>
        <status_str value="Command completed successfully"/>
```

| Field ID | Field output | Default position |
|---|---|---|
| **fqdn** | FQDN | 1 |
| **address** | Address | 2 |
| **base_dn** | Base DN | 3 |
| **has_certificate** | Has Certificate | 4 |
| **expiration_date** | Expiration Date | 5 |
| **valid_certificate** | Valid Certificate | N/A |
| **accessible** | Accessible | N/A |
| **port** | Port | 6 |
| **secure_port** | Secure Port | 7 |

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Disallowed |
| Read-only users | Allowed |
| Technicians | Disallowed |

## Listing LDAP server users

Use the **ldap_user_list** command to list LDAP server users.

```
ldap_user_list role=Category [ domain=DomainName ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **role** | Enumeration | The role of the users to be retrieved from the LDAP server. The available roles are: storageadmin and readonly. | Y | N/A |
| **domain** | Object name | The domain name. | N | All Domains |

This command retrieves a list of users from the LDAP server by a specific role.

| Field ID | Field output | Default position |
|---|---|---|
| **user_name** | User Name | 1 |
| **user_role** | Role | 2 |

**Example:**

```
ldap_user_list role=storageadmin
```

**Output:**

```
User Name       Role
--------------- -----------
readonly_user   Read Only
restldapread    Read Only
test_readonly   Read Only
xivreadonly     Read Only
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Allowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

### Return codes

- **LDAP_AUTHENTICATION_IS_NOT_ACTIVE**

  LDAP authentication is not active.
- **LDAP_IS_NOT_FULLY_CONFIGURED**

  LDAP is not fully configured.

  **Troubleshooting:** Check LDAP settings.
- **NO_LDAP_SERVERS_ARE_DEFINED**

  No LDAP servers are defined in the system
- **NO_LDAP_SERVERS_WITH_CERTIFICATE_ARE_DEFINED**

  No LDAP servers with an LDAP certificate are defined in the system.
- **LOGIN_FAILURE_XIV_USER_NOT_AUTHENTICATED_BY_LDAP_SERVER**

  XIV User '*XIV User*' was not authenticated by LDAP server '*Server FQDN*'.
- **LOGIN_FAILURE_LDAP_SERVER_UNREACHABLE**

  No LDAP server can be reached.
- **LOGIN_FAILURE_INVALID_BASE_DN**

  The base dn of server '*Server FQDN*' is invalid.

## Listing LDAP-based authentication mode

Use the **ldap_mode_get** command to list LDAP-based authentication mode.

```
ldap_mode_get
```

The command succeeds regardless of whether the LDAP server is accessible.

**Example:**

```
ldap_mode_get
```

**Output:**

```
Mode
Inactive
```

| Field ID | Field output | Default position |
|---|---|---|
| **mode** | Mode | 1 |

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Disallowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

## Enabling or disabling LDAP-based authentication mode

Use the **ldap_mode_set** command to enable or disable LDAP-based authentication mode.

```
ldap_mode_set  mode=Mode
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **mode** | Boolean | The required state of LDAP authentication. Available values: Active, Inactive. | Y |

**Example:**

```
ldap_mode_set mode=active
```

**Output:**

```
Command completed successfully.
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Disallowed |
| Read-only users | Disallowed |

| User Category | Permission |
|---|---|
| Technicians | Disallowed |

### Warnings

- **USERS_ASSOCIATED_TO_DOMAIN_ARE_YOU_SURE_YOU_WANT_TO_ENABLE_LDAP_AUTHENTICATION**

  There are users associated with domains. Are you sure you want to enable LDAP authentication?
- **ARE_YOU_SURE_YOU_WANT_TO_ENABLE_LDAP_AUTHENTICATION**

  Are you sure you want to enable LDAP authentication?
- **ARE_YOU_SURE_YOU_WANT_TO_DISABLE_LDAP_AUTHENTICATION**

  Are you sure you want to disable LDAP authentication?

### Return codes

- **LDAP_IS_NOT_FULLY_CONFIGURED**

  LDAP is not fully configured.

  **Troubleshooting:** Check LDAP settings.
- **NO_LDAP_SERVERS_WITH_CERTIFICATE_ARE_DEFINED**

  No LDAP servers with an LDAP certificate are defined in the system.
- **NO_LDAP_SERVERS_ARE_DEFINED**

  No LDAP servers are defined in the system

## Updating an LDAP server definition

Use the **ldap_update_server** command to update an existing server configuration.

```
ldap_update_server  fqdn=Fqdn [ address=Address ] [ base_dn=LdapDn ]  [ port=PortNum ]
[ secure_port=PortNum ]
   [ certificate=PemCertificate  | remove_certificate=<no|yes> ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **fqdn** | N/A | FQDN of the LDAP server. | Y | N/A |
| **address** | N/A | IP address of the LDAP server. | N | none |
| **certificate** | N/A | The public certificate or certificate chain of the LDAP server to be added (see below for details). | N | no certificate |
| **remove_certificate** | Boolean | Defines whether to remove the certificate. | N | no |
| **base_dn** | N/A | Base_DN of the LDAP directory. | N | none |
| **port** | Integer | The port number. | N | none |
| **secure_port** | Integer | The secure port number. | N | none |

**The certificate parameter:**

The value of the **certificate** parameter is the content of a PEM file with asterisks instead of newlines. Chained certificates are supported. The total maximal length of a PEM file holding chained certificates (leaf first, root last) is 15360 characters (including the asterisk characters). In Windows, you can drag-and-drop a PEM file from the Windows Explorer to the appropriate location in the XCLI session window; the content will be added automatically.

**Example:**

```
ldap_update_server fqdn=ldap.example.com address=192.0.2.1
        remove_certificate=yes
```

**Output:**

```
Command completed successfully.
```

## Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Disallowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

## Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_UPDATE_LDAP_SERVER**

  Are you sure you want to update the LDAP server configuration?

## Return codes

- **LDAP_SERVER_NOT_FOUND**

  No LDAP server with the specified FQDN is defined in the system.
- **ADDRESS_CURRENTLY_ASSOCIATED_WITH_ANOTHER_LDAP_SERVER**

  The specified IP address is currently associated with another LDAP server.
- **NO_UPDATE_PARAMETERS_SPECIFIED**

  No LDAP server parameters were specified for the update.
- **SSL_CERTIFICATE_CHAIN_EMPTY**

  No certificates were found in the input.
- **SSL_CERTIFICATE_HAS_EXPIRED**

  The SSL certificate has expired.
- **SSL_CERTIFICATE_INVALID_FORMAT**

  The SSL certificate format is invalid or corrupted.
- **SSL_CERTIFICATE_ISSUER_NOT_FOUND**

  The SSL certificate issuer was not found in the certificate chain.
- **SSL_CERTIFICATE_NOT_YET_VALID**

  The SSL certificate is not yet valid.

- **SSL_CERTIFICATE_VERIFICATION_FAILED**

  The SSL certificate chain verification failed.
- **SSL_CERTIFICATE_VERIFICATION_INTERNAL_ERROR**

  The SSL certificate verification has failed because of an internal system error.

## Removing an LDAP server definition

Use the **ldap_remove_server** command to remove an LDAP server definition.

```
ldap_remove_server  fqdn=Fqdn
```

### Parameters

| Name | Description | Mandatory |
|---|---|---|
| **fqdn** | FQDN of the server to remove. | Y |

**Example:**

```
ldap_remove_server fqdn=cloud.example.com
```

**Output:**

```
Command completed successfully.
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Disallowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

### Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_REMOVE_LDAP_SERVER**

  Are you sure you want to remove the LDAP server?

### Return codes

- **LDAP_SERVER_NOT_FOUND**

  No LDAP server with the specified FQDN is defined in the system.
- **LDAP_IS_ACTIVE_BUT_THIS_IS_THE_LAST_SERVER**

  Deleting the last LDAP server is illegal when LDAP authentication is active.
- **LDAP_USES_SSL_BUT_THIS_IS_THE_LAST_SERVER_WITH_CERTIFICATE**

  Deleting the last LDAP server with a valid SSL certificate is illegal when LDAP authentication is active and uses SSL.

# Launching the ldapsearch utility

Use the **ldap_search** command to launch the ldapsearch utility.

```
ldap_search  user=UserName fqdn=Fqdn [ second_cmd=<yes|no> password=Password ]
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **user** | Object name | The username to search for. | Y | N/A |
| **password** | N/A | The user password to search for. | N | empty |
| **second_cmd** | Boolean | Defines whether to invoke the second lsearch command. | N | no |
| **fqdn** | N/A | FQDN of LDAP server to query. | Y | N/A |

There are 2 LDAP search commands executed in the authentication process. The second one can be issued by setting the **second_cmd** parameter to yes.

**Example:**

```
ldap_search fqdn user password
```

**Output:**

```
Name           Index   Value
-------------- ------- ------------------------------------------------------
command_line   0       ldapsearch -H ldap://ldapwin2003.example.com:389...
returncode     0       0
stderr         0
stdout         0       dn: CN=employee,CN=Users,DC=xivldap2,DC=com
stdout         1       description: Group One
stdout         2       objectSid:: AQUAAAAAAAUVAAAAYcKhSnhmt01IPSuAbQQAAA==
stdout         3
stdout         4
```

| Field ID | Field output | Default position |
|---|---|---|
| **name** | Name | 1 |
| **index** | Index | 2 |
| **value** | Value | 3 |

## Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Disallowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

### Return codes

- **LOGIN_FAILURE_USER_NOT_AUTHENTICATED_BY_LDAP_SERVER**

  User *User Name* was not authenticated by LDAP server '*Server FQDN*'.
- **LOGIN_FAILURE_USER_MISSING_GROUP_ATTRIBUTE**

  User *User Name* is missing the group attribute '*Attribute*'.
- **LDAP_SERVER_NOT_FOUND**

  No LDAP server with the specified FQDN is defined in the system.
- **LOGIN_FAILURE_LDAP_SERVER_UNREACHABLE**

  No LDAP server can be reached.
- **LDAP_SERVER_NOT_DEFINED**

  LDAP server *Server FQDN* is not defined in the system.
- **LDAP_ROLE_UNRECOGNIZED**

  The LDAP role for the user is not recognized in the system.
- **LOGIN_FAILURE_USER_HAS_NO_RECOGNIZED_ROLE**

  User *User Name* has no recognized LDAP role.
- **LOGIN_FAILURE_USER_CANNOT_BE_UNIQUELY_AUTHENTICATED_BY_LDAP_SERVER**

  User *User Name* was not uniquely authenticated by LDAP server '*Server FQDN*'.
- **LOGIN_FAILURE_XIV_USER_NOT_AUTHENTICATED_BY_LDAP_SERVER**

  XIV User '*XIV User*' was not authenticated by LDAP server '*Server FQDN*'.
- **LOGIN_FAILURE_USER_HAS_MORE_THAN_ONE_RECOGNIZED_ROLE**

  User *User Name* has more than one recognized LDAP role.
- **LOGIN_FAILURE_USER_MISSING_ID_ATTRIBUTE**

  User *User Name* is missing the LDAP ID attribute '*Attribute*'.
- **USER_IS_PREDEFINED_IN_THE_SYSTEM**

  The user is predefined in the system.
- **LOGIN_FAILURE_INVALID_BASE_DN**

  The base dn of server '*Server FQDN*' is invalid.
- **LDAP_AUTHENTICATION_IS_NOT_ACTIVE**

  LDAP authentication is not active.

## Defining a new user

Use the **user_define** command to define a new user.

```
user_define user=UserName password=Password password_verify=Password category=Category
[ email_address=email ]
 [ area_code=AreaCode number=PhoneNumber ] [ domain=DomainList [ exclusive=<yes|no> ] ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **user** | Object name | User name. User names are lower case. | Y | N/A |

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **password** | N/A | Password of the user to be created. The password must have between 6 and 16 characters. Any symbols are allowed, except the following:<br>• double quotation (")<br>• single quotation or apostrophe (')<br>• grave accent (`)<br>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:<br>• storageadmin<br>• applicationadmin<br>• securityadmin<br>• readonly<br>• storageintegrationadmin | Y | N/A |
| **email_address** | N/A | Email address of this user. The email address specified here can be used for event notification. Entering this address is optional. The email address format is any legal email address. | N | N/A |
| **number** | N/A | Cellular phone number of the user for event notification via SMS, excluding the area code. Phone numbers and area codes can be a maximum of 63 digits, dashes (-) and periods (.) | N | N/A |
| **area_code** | N/A | Area code of the cellular phone number of the user. Phone numbers and area codes can be a maximum of 63 digits, dashes (-) and periods (.) | N | N/A |
| **domain** | N/A | The cluster will be attached to the specified domains. To specify more than one domain, separate them with a comma. To specify all existing domains, use "*". | N | none |

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **exclusive** | Boolean | Use yes to restrict the user to domain's objects. | N | yes |

Email address and phone number are optional and can be used for event notification. The category (user role) may be only one of those specified above (other categories contain only a single predefined user).

The maximum number of users is 128.

There are two system-wide predefined users:

| User name | User role |
|---|---|
| admin | Storage administrator |
| technician | Technician |

**Example:**

```
user_define user=xiv_user1 password=s0mePassw0rd password_verify=s0mePassw0rd
category=applicationadmin
```

**Output:**

```
Command completed successfully.
```

## Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Allowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

## Return codes

- **USER_NAME_ALREADY_EXISTS**

  The user name already exists.
- **MAX_USERS_REACHED**

  The number of defined users already reached the maximum.
- **PASSWORDS_DO_NOT_MATCH**

  Make sure that passwords are identical.
- **USER_PHONE_NUMBER_MUST_ACCOMPANY_AREA_CODE**

  The phone number must be indicated together with the area code.
- **LDAP_AUTHENTICATION_IS_ACTIVE**

  This command is not available while LDAP authentication is active.
- **DOMAIN_DOESNT_EXIST**

  The domain does not exist.
- **NON_EXCLUSIVE_USER_NOT_SECURITY_ADMIN**

Only a security administrator can add a user to a domain non-exclusively.

- **USER_CANNOT_BE_ADDED_TO_A_DOMAIN**

  The specified user cannot be associated with a domain.

- **SIA_MUST_BE_ASSOCIATED_WITH_A_DOMAIN**

  The storage integration administrator must be associated with a domain.

## Deleting a user

Use the **user_delete** command to delete a user.

```
user_delete user=UserName
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **user** | Object name | User to be deleted. | Y |

Existing objects created by this user will retain an empty user reference after the user has been deleted.

Two predefined users are set system-wide: Admin and Technician. Predefined users cannot be deleted or renamed.

**Example:**

```
user_delete user=user1
```

**Output:**

```
Command completed successfully.
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Allowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

### Return codes

- **LDAP_AUTHENTICATION_IS_ACTIVE**

  This command is not available while LDAP authentication is active.

- **USER_NAME_DOES_NOT_EXIST**

  The user name does not exist.

- **USER_CANNOT_BE_DELETED**

  The user cannot be deleted.

- **USER_IS_REFERRED_TO_BY_DEST**

  The user is referred to by an event destination and therefore cannot be deleted.

- **USER_OWNS_RECOVERY_KEY**

  The user owns a recovery key and therefore cannot be deleted or renamed.
- **REMOVAL_WOULD_CREATE_UNRESOLVABLE_REFERENCE_BETWEEN_USER_AND_USERGROUP**

  Completing this operation will result in a user referring to a user group that is not in its domain.

  **Troubleshooting:** Remove the reference explicitly and re-run the command.

## Adding users to user groups

Use the **user_group_add_user** command to add a user to a user group.

```
user_group_add_user user_group=UserGroup user=UserName
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **user_group** | Object name | User group into which the user is to be added. | Y |
| **user** | Object name | User to be added to the user group. | Y |

A user group can contain up to eight users.

A user may belong to only one user group.

Only users defined as Application Administrators can be assigned to a group.

This command fails when the user already belongs to the user group.

**Example:**

```
user_group_add_user user_group=ug1 user=user1
```

**Output:**

```
Command completed successfully.
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Allowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

### Return codes

- **USER_GROUP_NAME_DOES_NOT_EXIST**

  The user group name does not exist.
- **USER_NAME_DOES_NOT_EXIST**

  The user name does not exist.

- **USER_ALREADY_INCLUDED_IN_ANOTHER_GROUP**

  The user is included into another user group.
- **USER_GROUP_ALREADY_INCLUDES_USER**

  The user group already includes a user.
- **ONLY_APPLICATION_ADMIN_USERS_CAN_BE_GROUPED**

  User groups can only contain application administrators.
- **USER_GROUP_HAS_MAXIMUM_NUMBER_OF_USERS**

  The user group already contains the maximum number of users.
- **LDAP_AUTHENTICATION_IS_ACTIVE**

  This command is not available while LDAP authentication is active.

## Creating user groups

Use the **user_group_create** command to create a user group.

```
user_group_create user_group=UserGroup [ access_all=<yes|no> ]   [ ldap_role=LdapRole ]
[ domain=DomainList ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **user_group** | Object name | Name of the user group to be created. | Y | N/A |
| **access_all** | Boolean | Allows application administrators to perform their specified operations on all volumes and not just on a subset of the specific volumes. | N | no |
| **ldap_role** | String | The value representing the user group in LDAP. | N | [none] |
| **domain** | N/A | The user_group will be attached to the specified domains. To specify more than one domain, separate them with a comma. To specify all the existing domains, use "*". | N | none |

A user group is a group of application administrators who share the same set of snapshot creation limitations. After user groups are created, the limitations of all the users in a user group can be updated with a single command. These limitations are enforced by associating the user groups with hosts or clusters.

Storage administrators create user groups and control the various application administrator's permissions. Hosts and clusters can be associated with only a single user group. When a user belongs to a user group that is associated with a host, it is possible to manage snapshots of the volumes mapped to that host.

User groups have the following limitations:

- Only users who are defined as application administrators can be assigned to a group.
- A user can belong to only a single user group.
- A user group can contain up to eight users.

User and host associations have the following properties:

- User groups can be associated with both hosts and clusters. This allows limiting application administrator access to specific volumes.
- A host that is part of a cluster cannot also be associated with a user group.
- When a host is added to a cluster the host's associations are broken. Limitations on the management of volumes mapped to the host is controlled by the cluster's association.
- When a host is removed from a cluster, the host's associations become the cluster 's associations, this allows continued mapping of operations so that all scripts continue to work.

Application administrator access level:

- The **access_all** parameter can be specified for application administrators only. When it is specified, it means that the user has an application administrator access level to all volumes, and can perform operations on all volumes and not just on a subset of the specific volume.

**Example:**

```
user_group_create user_group=ug1 ldap_role="App Admin 1" access_all=yes
```

**Output:**

```
Command completed successfully.
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Disallowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

### Return codes

- **USER_GROUP_NAME_ALREADY_EXISTS**

  The user group name already exists.
- **MAX_USER_GROUPS_REACHED**

  The number of defined user groups already reached the maximum.
- **LDAP_ROLE_ALREADY_USED**

  This LDAP role is already in use in the LDAP configuration or in a user group.
- **DOMAIN_DOESNT_EXIST**

  The domain does not exist.

## Deleting a user group

Use the **user_group_delete** command to delete a user group.

```
user_group_delete user_group=UserGroup
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **user_group** | Object name | User group to be deleted. | Y |

A user group can be deleted, even when it is associated with hosts or clusters. It can be deleted while in LDAP Authentication mode.

A user group can be deleted, even when it contains users. Deleting the user group does not delete the users contained in this group.

**Example:**

```
user_group_delete user_group=ug1
```

**Output:**

```
Command completed successfully.
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Disallowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

### Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_DELETE_LDAP_USER_GROUP**

  One or more LDAP users might be associated with this user group. Are you sure you want to delete the user group?
- **ARE_YOU_SURE_YOU_WANT_TO_DELETE_USER_GROUP**

  Are you sure you want to delete the user group?
- **ARE_YOU_SURE_YOU_WANT_TO_DELETE_POPULATED_USER_GROUP**

  One or more internal users are associated with this user group. Are you sure you want to delete the user group?

### Return codes

- **USER_GROUP_NAME_DOES_NOT_EXIST**

  The user group name does not exist.

## Listing user groups

Use the **user_group_list** command to list all user groups or a specific one.

```
user_group_list [ user_group=UserGroup ] [ domain=DomainName ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **user_group** | Object name | The user group to be listed. | N | All user groups. |
| **domain** | Object name | The domain name. | N | All Domains |

All the users included in the user group are listed.

| Field ID | Field output | Default position |
|---|---|---|
| **name** | Name | 1 |
| **access_all** | Access All | 2 |
| **ldap_role** | LDAP Role | 3 |
| **users** | Users | 4 |
| **creator** | Creator | N/A |

### Example:

```
user_group_list
```

### Output:

```
Name        Access All   LDAP Role   Users
myug1       yes          Group1
myOtherUG   yes          OtherGroup
ug1         yes          App Admin 1
ug2         yes          App Admin 2
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Allowed |
| Security administrator | Disallowed |
| Read-only users | Allowed |
| Technicians | Disallowed |

## Removing a user from a user group

Use the **user_group_remove_user** command to remove a user from a user group.

```
user_group_remove_user user_group=UserGroup user=UserName
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **user_group** | Object name | User group. | Y |
| **user** | Object name | User to be removed. | Y |

This command fails when the user does not belong to the user group.

Deleting the user group's mapping is done by removing the role association. The user group itself is not deleted.

**Example:**

```
user_group_remove_user user_group=ug1 user=user1
```

**Output:**

```
Command completed successfully.
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Allowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

### Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_REMOVE_USER**

  Are you sure you want to remove the user from the user group?

### Return codes

- **USER_GROUP_NAME_DOES_NOT_EXIST**

  The user group name does not exist.
- **USER_NAME_DOES_NOT_EXIST**

  The user name does not exist.
- **USER_GROUP_DOES_NOT_INCLUDE_USER**

  The user group does not include any user.
- **LDAP_AUTHENTICATION_IS_ACTIVE**

  This command is not available while LDAP authentication is active.

## Renaming user groups

Use the **user_group_rename** command to rename a user group.

```
user_group_rename user_group=UserGroup new_name=Name
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **user_group** | Object name | User group to be renamed. | Y |
| **new_name** | Object name | New name of the user group. | Y |

**Example:**

```
user_group_rename user_group=ug1 new_name=ug2
```

**Output:**

```
Command completed successfully.
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Disallowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

### Return codes

- **USER_GROUP_NAME_DOES_NOT_EXIST**

  The user group name does not exist.
- **USER_GROUP_NAME_ALREADY_EXISTS**

  The user group name already exists.

## Updating a user group

Use the **user_group_update** command to update a user group.

```
user_group_update  user_group=UserGroup [ ldap_role=LdapRole ] [ access_all=<yes|no> ]
[ domain=DomainList ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **user_group** | Object name | The name of the user group to be updated. | Y | N/A |
| **ldap_role** | String | The value representing the user group in LDAP. | N | Keep current LDAP role. |
| **access_all** | Boolean | Assigns application administration access level for all volumes. | N | no |
| **domain** | N/A | The **user_group** will be attached to the specified domains. To specify more than one domain, separate them with a comma. To specify all the existing domains, use "*". | N | none |

**Example:**

```
user_group_update user_group=ug1 ldap_role="App Admin 1" access_all=yes
```

**Output:**

```
Command completed successfully.
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Disallowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

### Return codes

- **USER_GROUP_NAME_DOES_NOT_EXIST**

  The user group name does not exist.
- **LDAP_ROLE_ALREADY_USED**

  This LDAP role is already in use in the LDAP configuration or in a user group.
- **DOMAIN_DOESNT_EXIST**

  The domain does not exist.
- **REMOVAL_WOULD_CREATE_UNRESOLVABLE_REFERENCE_BETWEEN_HOST_AND_USERGROUP**

  Completing this operation will result in a host referring to a user group that is not in its domain.

  **Troubleshooting:** Remove the reference explicitly and re-run the command.
- **REMOVAL_WOULD_CREATE_UNRESOLVABLE_REFERENCE_BETWEEN_CLUSTER_AND_USERGROUP**

  Completing this operation will result in a cluster referring to a user group that is not in its domain.

  **Troubleshooting:** Remove the reference explicitly and re-run the command.
- **REMOVAL_WOULD_CREATE_UNRESOLVABLE_REFERENCE_BETWEEN_USER_AND_USERGROUP**

  Completing this operation will result in a user referring to a user group that is not in its domain.

  **Troubleshooting:** Remove the reference explicitly and re-run the command.

## Listing users

Use the **user_list** command to list all users or a specific user.

```
user_list [ user=UserName | show_users=<all|active> ] [ domain=DomainName ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **user** | Object name | The user to be listed. | N | All users. |
| **show_users** | Enumeration | Indicates whether all internal users will be listed, or only internal users that are active. | N | active |
| **domain** | Object name | The domain name. | N | All Domains |

The following information is listed:

- User name: Lower case
- Category
- Email address
- Phone number
- Phone area code
- Containing user group

Passwords are not shown in the list.

| Field ID | Field output | Default position |
|---|---|---|
| **name** | Name | 1 |
| **category** | Category | 2 |
| **group** | Group | 3 |
| **active** | Active | 4 |
| **email_address** | Email Address | 5 |
| **area_code** | Area Code | 6 |
| **number** | Phone Number | 7 |
| **access_all** | Access All | 8 |
| **id** | ID | N/A |
| **creator** | Creator | N/A |
| **creator_category** | Creator Category | N/A |

**Example:**

```
user_list
```

**Output:**

```
Name              Category          Group
xiv_development   xiv_development
yes
xiv_maintenance   xiv_maintenance
yes
admin             storageadmin
yes
technician        technician               yes
```

## Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Allowed |
| Security administrator | Allowed |
| Read-only users | Allowed |
| Technicians | Disallowed |

# Renaming users

Use the **user_rename** command to rename a user.

```
user_rename user=UserName new_name=Name
```

## Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **user** | Object name | The user to be renamed. User names are lowercase. | Y |
| **new_name** | Object name | New name of the user. | Y |

This command renames a user.

**Example:**

```
user_rename user=admin new_name=storage_admin
```

**Output:**

```
Command completed successfully.
```

## Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Allowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

## Return codes

- **USER_NAME_DOES_NOT_EXIST**

  The user name does not exist.
- **USER_NAME_ALREADY_EXISTS**

  The user name already exists.
- **USER_CANNOT_BE_RENAMED**

  The user cannot be renamed.
- **LDAP_AUTHENTICATION_IS_ACTIVE**

  This command is not available while LDAP authentication is active.
- **USER_OWNS_RECOVERY_KEY**

  The user owns a recovery key and therefore cannot be deleted or renamed.
- **OPERATION_NOT_ALLOWED_ON_DESIGNATED_MSM_USER**

  The designated MSM user cannot be renamed and cannot be exclusively associated with a domain.

# Updating a user definition

Use the **user_update** command to update a user definition.

```
user_update user=UserName [ password=Password password_verify=Password ] [ email_address=email ]
 [ area_code=AreaCode ] [ number=PhoneNumber ] [ exclusive=<yes|no> ]
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **user** | Object name | The name of the user to be updated. User names are lower case. | Y | N/A |
| **password** | N/A | New password. Users can only change their own passwords. Password of the user to be updated. The password must have between 6 and 16 characters. Any symbols are allowed, except the following:<br>• double quotation (")<br>• single quotation or apostrophe (')<br>• grave accent (`)<br>Passwords are case sensitive. | N | Retains the current password. |
| **password_verify** | N/A | Verification of the password: Must be equal to the password. | N | Retains the current password. |
| **email_address** | N/A | Email address of the user (for event notification). | N | Leaves the current email address. |
| **number** | N/A | Cellular phone number of the user (for event notification via SMS) excluding the area code. | N | Leaves the current number. |
| **area_code** | N/A | Area code of the cellular phone number of the user. | N | Leaves the current area code. |
| **exclusive** | Boolean | This parameter can be set only by security administrator. If set to "yes", the user will be removed from the global domain. If set to "no", the user will get permissions on the global domain. | N | Leaves the current value. |

Users can change their own user parameters, such as password, email address, area code, and phone number. **Only** Storage administrators and Security administrators can change parameters of other users. In addition, a security administrator can change the exclusive parameter for all users.

When updating a user password, the following restrictions apply:

| User type | Predefined name (if exists) | Can change own password? | Can change other user passwords? | Limitations |
|---|---|---|---|---|
| Storage administrator | admin | Y | Y | Cannot change the password of the predefined Technician user |
| Storage administrator | (customer defined) | Y | N | N/A |
| Security administrator | (customer defined) | Y | Y | Cannot change the password of:<br>• other Security administrators<br>• the predefined Technician user<br>• the predefined Storage administrator 'admin' user |
| Technician | technician | N | N | N/A |
| Other (customer defined) | (customer defined) | Y | N | N/A |

**Example:**

```
user_update user=admin password=Passw0rd password_verify=Passw0rd
```

**Output:**

```
Command completed successfully.
```

## Access control

| User Category | Permission | Condition |
|---|---|---|
| Storage administrator | Conditionally Allowed | A user other than admin may only change its own configuration. |
| Storage integration administrator | Disallowed | N/A |
| Application administrator | Conditionally Allowed | A user of this category may only change its own configuration. |
| Security administrator | Conditionally Allowed | A user of this category may only change its own configuration. |
| Read-only users | Conditionally Allowed | A user other than admin may only change its own configuration. |
| Technicians | Disallowed | N/A |

## Return codes

- **USER_NAME_DOES_NOT_EXIST**

  The user name does not exist.
- **PASSWORDS_DO_NOT_MATCH**

  Make sure that passwords are identical.
- **USER_PHONE_NUMBER_MUST_ACCOMPANY_AREA_CODE**

  The phone number must be indicated together with the area code.
- **ADMIN_CAN_NOT_CHANGE_TECHNICIAN_USER**

  The administrator is not allowed to modify the details of a technician.

- **SMS_DESTINATION_REFERS_TO_USER**

  An SMS destination refers to the user and therefore must be defined by a phone number and an area code.
- **EMAIL_DESTINATION_REFERS_TO_USER**

  An email destination refers to the user, and therefore must be defined by an email address.
- **USER_NOT_ALLOWED_TO_CHANGE_OTHER_USERS**

  This user is not allowed to modify the details of other users.
- **USER_NOT_ALLOWED_TO_HAVE_PHONE_NUMBER**

  This user is not allowed to have a phone number.
- **USER_NOT_ALLOWED_TO_HAVE_EMAIL_ADDRESS**

  This user is not allowed to have an email address.
- **USER_NOT_ALLOWED_TO_CHANGE_PASSWORDS**

  This user cannot change passwords of other users.
- **USER_CANNOT_BE_UPDATED_WHILE_LDAP_AUTHENTICATION_IS_ACTIVE**

  The user cannot be updated while LDAP authentication is active.
- **NON_EXCLUSIVE_USER_NOT_SECURITY_ADMIN**

  Only a security administrator can add a user to a domain non-exclusively.
- **PREDEFINED_USER_CANNOT_BE_RESTRICTED_TO_DOMAIN**

  The specified user cannot be exclusively associated with a domain.
- **SIA_MUST_BE_ASSOCIATED_WITH_A_DOMAIN**

  The storage integration administrator must be associated with a domain.
- **USER_IS_REFERRED_TO_BY_DEST**

  The user is referred to by an event destination and therefore cannot be deleted.
- **OPERATION_NOT_ALLOWED_ON_DESIGNATED_MSM_USER**

  The designated MSM user cannot be renamed and cannot be exclusively associated with a domain.
- **REMOVAL_WOULD_CREATE_UNRESOLVABLE_REFERENCE_BETWEEN_USER_AND_USERGROUP**

  Completing this operation will result in a user referring to a user group that is not in its domain.

  **Troubleshooting:** Remove the reference explicitly and re-run the command.

## Creating a new domain

Use the **domain_create** command to create a domain.

```
domain_create domain=DomainName [ size=GB ] [ max_pools=MaxPools ] [ max_volumes=MaxVolumes ]
 [ max_cgs=MaxCGs ] [ max_mirrors=MaxMirrors ] [ max_dms=MaxDataMigrations ]
 [ perf_class=perfClassName ] [ ldap_id=LdapRole ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **domain** | Object name | The name of the domain to be created. | Y | N/A |
| **size** | Integer | Defines the sum of the sizes of all the pools associated with the domain, in gigabytes. | N | 0 |

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **max_pools** | Positive integer | The maximum number of pools that can be associated with this domain. | N | 0 |
| **max_volumes** | Positive integer | The maximum number of volumes that can be associated with all the pools in this domain. | N | 0 |
| **max_cgs** | Integer | The maximum number of consistency groups that can be associated with this domain. | N | 512 |
| **max_mirrors** | Positive integer | The maximum number of mirrors that can be associated with this domain. | N | 0 |
| **max_dms** | Positive integer | The maximum number of data migrations that can be associated with this domain. | N | 0 |
| **perf_class** | Object name | Name of a performance class. | N | none |
| **ldap_id** | String | The name to be associated with this domain in LDAP. | N | The domain name |

**Example:**

```
domain_create domain=d1 size=1000
```

**Output:**

```
Command completed successfully.
```

## Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Disallowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

## Warnings

- **DO_YOU_WANT_TO_SHARE_LDAP_ID**

  The specified LDAP ID is already in use. Are you sure you want to share the same LDAP ID?
- **DOMAIN_SIZE_SMALL**

  The domain size is smaller than the minimal pool size. Are you sure?

### Return codes

- **DOMAIN_ALREADY_EXISTS**

  A domain with this name already exists.
- **DOMAIN_MAX_REACHED**

  The maximum allowed number of domain objects is already reached.
- **PERF_CLASS_BAD_NAME**

  The performance class does not exist.
- **PERF_CLASS_ASSOCIATED_WITH_HOSTS**

  Performance class *Performance Class* is already being used by a host.
- **PERF_CLASS_ASSOCIATED_WITH_VOLUMES**

  Performance class *Performance Class* is already being used by a volume.
- **DOMAIN_INSUFFICIENT_CAPACITY**

  There is not enough capacity available for the domain.
- **DOMAIN_INSUFFICIENT_VOLUMES**

  There are not enough volumes available for the domain.
- **DOMAIN_INSUFFICIENT_POOLS**

  There are not enough pools available for the domain.
- **DOMAIN_INSUFFICIENT_CGS**

  There are not enough consistency groups available for the domain.
- **DOMAIN_INSUFFICIENT_MIRRORS**

  There are not enough mirrors available for the domain.
- **DOMAIN_INSUFFICIENT_DMS**

  There are not enough data migrations available for the domain.

## Updating a domain definition

Use the **domain_update** command to update a domain definition.

```
domain_update domain=DomainName [ size=GB ] [ max_pools=MaxPools ] [ max_volumes=MaxVolumes ]
 [ max_cgs=MaxCGs ] [ max_mirrors=MaxMirrors ] [ max_dms=MaxDataMigrations ]
 [ perf_class=perfClassName ] [ ldap_id=LdapRole ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **domain** | Object name | The name of the domain to be updated. | Y | N/A |
| **size** | Integer | Defines the sum of the sizes of all the pools associated with the domain, in gigabytes. | N | Current value. |
| **max_pools** | Positive integer | The maximum number of pools that can be associated with this domain. | N | Current value. |

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **max_volumes** | Positive integer | The maximum number of volumes that can be associated with all the pools in this domain. | N | Current value. |
| **max_cgs** | Integer | The maximum number of consistency groups that can be associated with this domain. | N | Current value. |
| **max_mirrors** | Positive integer | The maximum number of mirrors that can be associated with this domain. | N | Current value. |
| **max_dms** | Positive integer | The maximum number of data migrations that can be associated with this domain. | N | Current value. |
| **perf_class** | Object name | Name of a performance class. | N | Current value. |
| **ldap_id** | String | The name to be associated with this domain in LDAP. | N | Current value. |

**Example:**

```
domain_update domain=d1 size=10000 max_pools=5 max_volumes=100
```

**Output:**

```
Command completed successfully.
```

## Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Disallowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

## Warnings

- **DO_YOU_WANT_TO_SHARE_LDAP_ID**

  The specified LDAP ID is already in use. Are you sure you want to share the same LDAP ID?
- **DOMAIN_SIZE_SMALL**

  The domain size is smaller than the minimal pool size. Are you sure?

## Return codes

- **DOMAIN_DOESNT_EXIST**

  The domain does not exist.
- **PERF_CLASS_BAD_NAME**

The performance class does not exist.

- **PERF_CLASS_ASSOCIATED_WITH_HOSTS**

  Performance class *Performance Class* is already being used by a host.
- **PERF_CLASS_ASSOCIATED_WITH_VOLUMES**

  Performance class *Performance Class* is already being used by a volume.
- **DOMAIN_SIZE_TOO_SMALL**

  The domain usage exceeds the requested size.
- **DOMAIN_MAX_VOLUMES_TOO_SMALL**

  The actual number of volumes in the domain already exceeds the specified number.
- **DOMAIN_MAX_MIRRORS_TOO_SMALL**

  The actual number of mirrors in the domain already exceeds the specified number.
- **DOMAIN_MAX_DMS_TOO_SMALL**

  The actual number of data migrations in the domain already exceeds the specified number.
- **DOMAIN_MAX_CGS_TOO_SMALL**

  The actual number of consistency groups in the domain already exceeds the specified number.
- **DOMAIN_MAX_POOLS_TOO_SMALL**

  The actual number of pools in the domain already exceeds the specified number.
- **DOMAIN_INSUFFICIENT_CAPACITY**

  There is not enough capacity available for the domain.
- **DOMAIN_INSUFFICIENT_VOLUMES**

  There are not enough volumes available for the domain.
- **DOMAIN_INSUFFICIENT_POOLS**

  There are not enough pools available for the domain.
- **DOMAIN_INSUFFICIENT_MIRRORS**

  There are not enough mirrors available for the domain.
- **DOMAIN_INSUFFICIENT_CGS**

  There are not enough consistency groups available for the domain.
- **DOMAIN_INSUFFICIENT_DMS**

  There are not enough data migrations available for the domain.

## Renaming a domain

Use the **domain_rename** command to rename a domain.

```
domain_rename domain=DomainName new_name=Name
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **new_name** | Object name | Name of the domain. | Y |
| **domain** | Object name | New name of the domain. | Y |

**Example:**

```
domain_rename domain=domain1 new_name=domain2
```

**Output:**

```
Command completed successfully
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Disallowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

### Return codes

- **DOMAIN_ALREADY_EXISTS**

  A domain with this name already exists.
- **DOMAIN_DOESNT_EXIST**

  The domain does not exist.

## Deleting a domain

Use the **domain_delete** command to delete a domain.

```
domain_delete domain=DomainName
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **domain** | Object name | The name of the domain to delete. | Y |

**Example:**

```
domain_delete domain=domain1
```

**Output:**

```
Command completed successfully
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Disallowed |

| User Category | Permission |
|---|---|
| Read-only users | Disallowed |
| Technicians | Disallowed |

### Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_DELETE_DOMAIN**

  Are you sure you want to delete the domain *Domain*?

### Return codes

- **DOMAIN_DOESNT_EXIST**

  The domain does not exist.
- **DOMAIN_HAS_POOL**

  One or more pools are still in the domain.
- **DOMAIN_HAS_USER_GROUP**

  One or more user groups are associated with this domain.
- **DOMAIN_HAS_USER**

  One or more users are associated with this domain.
- **DOMAIN_HAS_SCHEDULE**

  One or more schedules are associated with this domain.
- **DOMAIN_HAS_DEST**

  One or more destinations are associated with this domain.
- **DOMAIN_HAS_DESTGROUP**

  One or more destination groups are associated with this domain.

## Listing domains

List all domains or the specified one.

```
domain_list [ domain=DomainName ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **domain** | Object name | Name of a domain. | N | All domains. |

When the **domain** parameter is provided, only the specified domain is listed.

**Example:**

```
domain_list domain=d1
```

**Output:**

```
Tabular output

Name       DN         Soft   Free Soft   Hard   Free Hard
Domain1    Domain1    1703   0           1703   0
Domain2    Domain2    1703   1703        1703   1703
Domain3    Domain3    1703   1600        1703   1600
Domain4    Domain4    1703   1703        1703   1703
Domain5    Domain5    1703   1703        1703   1703

XML output
<domain id="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**

  The domain does not exist.

## Listing objects in domains

Use the **domain_list_objects** command to list objects attached to domain(s).

```
domain_list_objects  [ domain=DomainName ] [ type=ObjectType [ name=ObjectName ] ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **domain** | Object name | Name of a domain. | N | All domains. |
| **type** | Enumeration | The object type to list: target, host, cluster, schedule, usergroup, dest, destgroup or rule. | N | All object types. |
| **name** | Object name | Name of an object. | N | All object names. |

This command is used for listing objects in the system per domain.

**Example:**

```
domain_list_objects domain=d1
```

**Output:**

```
Domain   Type       Object
-------- ---------- --------------
d1       cluster    c1
d1       host       MyHost
d1       schedule   min_interval
d1       schedule   never
```

| Field ID | Field output | Default position |
|---|---|---|
| **domain_name** | Domain | 1 |
| **object_type** | Type | 2 |
| **object_name** | Object | 3 |

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Allowed |
| Application administrator | Allowed |
| Security administrator | Allowed |
| Read-only users | Allowed |
| Technicians | Allowed |

### Return codes

- **DOMAIN_DOESNT_EXIST**

  The domain does not exist.

- **TOO_MANY_OBJECTS**

  There are too many objects to output. Re-run the command by using the command's parameters to filter the output by domain or object type.

## Listing the global domain

Use the **domain_global_list** to list the global domain.

```
domain_global_list
```

**Example:**

```
domain_global_list
```

| Field ID | Field output | Default position |
|---|---|---|
| **name** | Name | 1 |
| **ldap_id** | LDAP ID | 2 |
| **size** | Size | 3 |
| **size_MiB** | Size (MiB) | N/A |
| **total_pool_size** | Total Pools (GB) | 4 |
| **total_pool_size_MiB** | Total Pools (MiB) | N/A |
| **empty_space** | Empty (GB) | 5 |
| **empty_space_MiB** | Empty (MiB) | N/A |
| **max_pools** | Max Pools | 6 |
| **used_pools** | Pools | 7 |
| **max_volumes** | Max Volumes | 8 |
| **used_volumes** | Volumes | 9 |
| **max_mirrors** | Max Mirrors | 10 |
| **used_mirrors** | Mirrors | 11 |
| **max_dms** | Max Data Migrations | 12 |
| **used_dms** | Data Migrations | 13 |
| **max_cgs** | Max CGs | 14 |

| Field ID | Field output | Default position |
|---|---|---|
| **used_cgs** | CGs | 15 |
| **perf_class** | Performance Class | 16 |
| **managed** | Managed | 17 |
| **id** | ID | N/A |

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Disallowed |
| Read-only users | Allowed |
| Technicians | Disallowed |

## Attaching an object to a domain

Use the **domain_attach_object** command to associate an object with a domain.

```
domain_attach_object domain=DomainName type=ObjectType name=ObjectName
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **domain** | Object name | The name of the domain. | Y |
| **type** | Enumeration | The object type to attach to the domain. It can be: target, host, cluster, schedule, usergroup, dest, destgroup or rule. | Y |
| **name** | Object name | The object name. | Y |

**Example:**

```
domain_attach_object domain=d1 type=host name=MyHost
```

**Output:**

```
Command completed successfully.
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Disallowed |
| Read-only users | Disallowed |

| User Category | Permission |
|---|---|
| Technicians | Disallowed |

### Return codes

- **DOMAIN_DOESNT_EXIST**

  The domain does not exist.
- **TARGET_BAD_NAME**

  The target name does not exist.
- **HOST_BAD_NAME**

  The host name does not exist.
- **CLUSTER_BAD_NAME**

  The cluster name does not exist.
- **USER_GROUP_NAME_DOES_NOT_EXIST**

  The user group name does not exist.
- **SCHEDULE_DOES_NOT_EXIST**

  The specified schedule does not exist.
- **DEST_NAME_DOES_NOT_EXIST**

  The destination name does not exist.
- **DESTGROUP_NAME_DOES_NOT_EXIST**

  The destination group name does not exist.
- **EVENT_RULE_NAME_DOES_NOT_EXIST**

  The event rule name does not exist.
- **USER_IS_NOT_IN_DESTINATION_DOMAINS**

  The user must be included in destination domains.
- **DESTINATION_IS_NOT_IN_DESTGROUP_DOMAINS**

  The destination must be included in the destination group domains.
- **DESTINATION_IS_NOT_IN_RULE_DOMAINS**

  The destination must be included in rule domains.
- **DESTGROUP_IS_NOT_IN_RULE_DOMAINS**

  The destination groups must be included in rule domains.
- **ESCALATION_RULE_NOT_IN_RULE_DOMAINS**

  An escalation rule must belong to rule domains.
- **CLUSTER_HAS_HOSTS_UNASSOCIATED_WITH_DOMAIN**

  The cluster cannot be attached, because it includes hosts that are not in the specified domain.
- **RESOURCE_ALREADY_ASSOCIATED_WITH_THIS_DOMAIN**

  The resource is already associated with this domain.

## Disassociating object from a domain

Use the **domain_detach_object** command to disassociate object from a domain.

```
domain_detach_object domain=DomainName type=ObjectType name=ObjectName
```

## Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **domain** | Object name | The name of the domain. | Y |
| **type** | Enumeration | The object type to disassociate from the domain. It can be: target, host, cluster, schedule, usergroup, dest, destgroup, or rule. | Y |
| **name** | Object name | The object name. | Y |

The object is disassociated from mapped or bound objects that belong to the domain.

**Example:**

```
domain_detach_object domain=d1 type=host name=MyHost
```

**Output:**

```
Command completed successfully.
```

## Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Disallowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

## Return codes

- **DOMAIN_DOESNT_EXIST**

  The domain does not exist.
- **RESOURCE_NOT_ASSOCIATED_WITH_THIS_DOMAIN**

  The resource is not associated with this domain.
- **DOMAIN_VOLUME_MAPPED_TO_HOST**

  A domain volume is mapped to this host.
- **DOMAIN_VOLUME_MAPPED_TO_CLUSTER**

  Cluster has a volume in the domain mapped to it.
- **REMOVAL_WOULD_CREATE_UNRESOLVABLE_REFERENCE_BETWEEN_HOST_AND_USERGROUP**

  Completing this operation will result in a host referring to a user group that is not in its domain.

  **Troubleshooting:** Remove the reference explicitly and re-run the command.
- **REMOVAL_WOULD_CREATE_UNRESOLVABLE_REFERENCE_BETWEEN_CLUSTER_AND_USERGROUP**

  Completing this operation will result in a cluster referring to a user group that is not in its domain.

  **Troubleshooting:** Remove the reference explicitly and re-run the command.
- **HOST_PART_OF_ATTACHED_CLUSTER**

The host is part of a cluster and cannot be handled individually.

- **DOMAIN_TARGET_IN_USE**

  The target domain cannot be removed, because it includes a volume used in a mirror, data migration, or OLVM relationship.

- **DOMAIN_VOLUME_BOUND_TO_HOST**

  A domain volume is bound to this host via an ALU.

- **TARGET_BAD_NAME**

  The target name does not exist.

- **HOST_BAD_NAME**

  The host name does not exist.

- **CLUSTER_BAD_NAME**

  The cluster name does not exist.

- **USER_GROUP_NAME_DOES_NOT_EXIST**

  The user group name does not exist.

- **SCHEDULE_DOES_NOT_EXIST**

  The specified schedule does not exist.

- **DEST_NAME_DOES_NOT_EXIST**

  The destination name does not exist.

- **DESTGROUP_NAME_DOES_NOT_EXIST**

  The destination group name does not exist.

- **EVENT_RULE_NAME_DOES_NOT_EXIST**

  The event rule name does not exist.

- **DETACH_WOULD_MAKE_OBJECT_INACCESSIBLE**

  Detaching the object will render it unassociated with any domain, and therefore inaccessible.

  **Troubleshooting:** Delete the object, if it is no longer needed.

- **USER_IS_NOT_IN_DESTINATION_DOMAINS**

  The user must be included in destination domains.

- **DESTINATION_IS_NOT_IN_RULE_DOMAINS**

  The destination must be included in rule domains.

- **DESTINATION_IS_NOT_IN_DESTGROUP_DOMAINS**

  The destination must be included in the destination group domains.

- **REMOVAL_WOULD_CREATE_UNRESOLVABLE_REFERENCE_BETWEEN_USER_AND_USERGROUP**

  Completing this operation will result in a user referring to a user group that is not in its domain.

  **Troubleshooting:** Remove the reference explicitly and re-run the command.

- **DESTGROUP_IS_NOT_IN_RULE_DOMAINS**

  The destination groups must be included in rule domains.

- **ESCALATION_RULE_NOT_IN_RULE_DOMAINS**

  An escalation rule must belong to rule domains.

- **DOMAIN_SCHEDULE_IN_USE**

  The schedule is in use, and therefore cannot be moved to another domain.

- **DOMAIN_PROXY_VOLUME_MAPPED_TO_HOST**

A proxy domain volume is mapped to this host.

- **DOMAIN_PROXY_VOLUME_MAPPED_TO_CLUSTER**

  A proxy domain volume is mapped to this cluster.

## Associating users to a domain

Use the **domain_add_user** command to associate a user to a domain.

```
domain_add_user domain=DomainName user=UserName [ exclusive=<yes|no> ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **domain** | Object name | The name of the domain. | Y | N/A |
| **user** | Object name | The name of the user. | Y | N/A |
| **exclusive** | Boolean | Set to Yes to restrict the user to domain's objects. | N | yes |

**Example:**

```
domain_add_user domain=d1 user=d1_admin
```

**Output:**

```
Command completed successfully.
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Allowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

### Return codes

- **LDAP_AUTHENTICATION_IS_ACTIVE**

  This command is not available while LDAP authentication is active.
- **DOMAIN_DOESNT_EXIST**

  The domain does not exist.
- **USER_NAME_DOES_NOT_EXIST**

  The user name does not exist.
- **NON_EXCLUSIVE_USER_NOT_SECURITY_ADMIN**

  Only a security administrator can add a user to a domain non-exclusively.
- **USER_CANNOT_BE_ADDED_TO_A_DOMAIN**

The specified user cannot be associated with a domain.

- **PREDEFINED_USER_CANNOT_BE_RESTRICTED_TO_DOMAIN**

  The specified user cannot be exclusively associated with a domain.
- **DOMAIN_USER_EXIST**

  This user is already added to the domain.
- **USER_IS_REFERRED_TO_BY_DEST**

  The user is referred to by an event destination and therefore cannot be deleted.
- **OPERATION_NOT_ALLOWED_ON_DESIGNATED_MSM_USER**

  The designated MSM user cannot be renamed and cannot be exclusively associated with a domain.
- **SIA_MUST_BE_ASSOCIATED_WITH_A_DOMAIN**

  The storage integration administrator must be associated with a domain.
- **REMOVAL_WOULD_CREATE_UNRESOLVABLE_REFERENCE_BETWEEN_USER_AND_USERGROUP**

  Completing this operation will result in a user referring to a user group that is not in its domain.

  **Troubleshooting:** Remove the reference explicitly and re-run the command.

# Removing a user from a domain

Use the **domain_remove_user** command to remove a user from a domain.

```
domain_remove_user domain=DomainName user=UserName
```

## Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **domain** | Object name | The name of the domain. | Y |
| **user** | Object name | The name of the user. | Y |

**Example:**

```
domain_remove_user domain=d1 user=d1_admin
```

**Output:**

```
Command completed successfully.
```

## Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Allowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

### Warnings

- **USER_WILL_BE_DELETED**

  The user is associated with this domain only. This will delete the user from the system. Are you sure?

### Return codes

- **USER_NAME_DOES_NOT_EXIST**

  The user name does not exist.
- **DOMAIN_DOESNT_HAVE_THE_USER**

  The user is not attached to this domain.
- **DOMAIN_DOESNT_EXIST**

  The domain does not exist.
- **LDAP_AUTHENTICATION_IS_ACTIVE**

  This command is not available while LDAP authentication is active.
- **REMOVAL_WOULD_CREATE_UNRESOLVABLE_REFERENCE_BETWEEN_USER_AND_USERGROUP**

  Completing this operation will result in a user referring to a user group that is not in its domain.

  **Troubleshooting:** Remove the reference explicitly and re-run the command.
- **DOMAIN_USER_CANNOT_REMOVE_HIMSELF**

  Users cannot remove themselves from a domain.
- **USER_IS_REFERRED_TO_BY_DEST**

  The user is referred to by an event destination and therefore cannot be deleted.

## Adding a pool to a domain

Use the **domain_add_pool** command to add a pool to a domain.

```
domain_add_pool domain=DomainName pool=PoolName [ adjust=<yes|no> ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **domain** | Object name | The name of the domain. | Y | N/A |
| **pool** | Object name | The pool name. | Y | N/A |
| **adjust** | Boolean | Adjust domain resources. If 'adjust' is set to true, the resources of the global domain and destination domain are adjusted to accommodate the pool being moved. | N | no |

**Example:**

```
domain_add_pool domain=d1 pool=p1
```

**Output:**

```
Command completed successfully.
```

## Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Disallowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

## Return codes

- **POOL_DOES_NOT_EXIST**

  The storage pool does not exist.
- **NO_FREE_CAPACITY_IN_DOMAIN**

  There is not enough free space in the domain.
- **DOMAIN_DOESNT_EXIST**

  The domain does not exist.
- **POOL_ALREADY_ASSOCIATED_WITH_A_DOMAIN**

  The pool is already associated with a domain.
- **DOMAIN_MAX_POOLS_REACHED**

  The maximum allowed number of domain pools is already reached.
- **DOMAIN_MAX_MIRRORS_REACHED**

  The domain exceeds the maximum allowed number of mirrors.
- **DOMAIN_MAX_DMS_REACHED**

  The domain exceeds the maximum allowed number of data migrations.
- **DOMAIN_MAX_CONS_GROUPS_REACHED**

  The domain exceeds the maximum allowed number of consistency groups.
- **DOMAIN_MAX_VOLUMES_REACHED**

  The domain exceeds the maximum allowed number of volumes.
- **DOMAIN_USED_TARGET_NOT_IN_DESTINATION**

  A target that is used by mirror in the pool is not associated with the target domain.
- **DOMAIN_USED_SCHEDULE_NOT_IN_DESTINATION**

  A schedule used by a mirror in the pool is not associated with the target domain.
- **MAPPED_HOSTS_NOT_IN_DESTINATION**

  A host that is mapped to a volume in the pool is not associated with the target domain.
- **MAPPED_CLUSTERS_NOT_IN_DESTINATION**

  A cluster that is mapped to a volume in the pool is not associated with the target domain.
- **NO_SPACE**

  The system does not have enough free space for the requested storage pool size.

# Removing a pool from a domain

Use the **domain_remove_pool** command to remove a pool from a domain.

```
domain_remove_pool domain=DomainName pool=PoolName [ adjust=<yes|no> ]
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **domain** | Object name | The name of the domain. | Y | N/A |
| **pool** | Object name | The pool name. | Y | N/A |
| **adjust** | Boolean | Adjust domain resources. If set to True, the resources of the global domain and destination domain are adjusted to accommodate the pool being moved. | N | no |

**Example:**

```
domain_remove_pool domain=d1 pool=p1
```

**Output:**

```
Command completed successfully.
```

## Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Disallowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

## Return codes

- **DOMAIN_DOESNT_EXIST**

  The domain does not exist.
- **POOL_DOES_NOT_EXIST**

  The storage pool does not exist.
- **DOMAIN_DOESNT_HAVE_THE_POOL**

  The pool is not attached to this domain.
- **MAX_POOLS_REACHED**

  The maximum allowed number of storage pools is already reached.
- **MAX_MIRRORS_REACHED**

The maximum number of mirrors is already reached.

- **MAX_CONS_GROUPS_REACHED**

  The maximum allowed number of consistency groups is already reached.
- **MAX_VOLUMES_REACHED**

  The maximum allowed number of volumes is already reached.
- **MAX_DMS_REACHED**

  The maximum number of remote volumes (mirror/migration) is already reached.

  **Troubleshooting:** Delete unnecessary data migration objects.
- **NO_SPACE**

  The system does not have enough free space for the requested storage pool size.
- **DOMAIN_USED_SCHEDULE_NOT_IN_DESTINATION**

  A schedule used by a mirror in the pool is not associated with the target domain.

# Moving a pool from one domain to another

Use the **domain_move_pool** command to move a pool from one domain to another.

```
domain_move_pool pool=PoolName src_domain=DomainName dst_domain=DomainName [ adjust=<yes|no> ]
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **pool** | Object name | The name of the pool to be moved. | Y | N/A |
| **src_domain** | Object name | The source domain name. | Y | N/A |
| **dst_domain** | Object name | The destination domain name. | Y | N/A |
| **adjust** | Boolean | Adjust domain resources. If set to Yes, the resources of the domains are adjusted to accommodate the pool being moved. | N | no |

**Example:**

```
domain_move_pool pool=p1 src_domain=d1 dst_domain=d2
```

**Output:**

```
Command completed successfully.
```

## Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Disallowed |

| User Category | Permission |
|---|---|
| Read-only users | Disallowed |
| Technicians | Disallowed |

## Return codes

- **SOURCE_DOMAIN_DOES_NOT_EXIST**

  The source domain does not exist.
- **DESTINATION_DOMAIN_DOES_NOT_EXIST**

  The destination domain does not exist.
- **POOL_DOES_NOT_EXIST**

  The storage pool does not exist.
- **POOL_NOT_ASSOCIATED_WITH_SOURCE_DOMAIN**

  The pool is not associated with the source domain.
- **DOMAIN_MAX_POOLS_REACHED**

  The maximum allowed number of domain pools is already reached.
- **MAPPED_HOSTS_NOT_IN_DESTINATION**

  A host that is mapped to a volume in the pool is not associated with the target domain.
- **MAPPED_CLUSTERS_NOT_IN_DESTINATION**

  A cluster that is mapped to a volume in the pool is not associated with the target domain.
- **NO_FREE_CAPACITY_IN_DOMAIN**

  There is not enough free space in the domain.
- **DOMAIN_USED_SCHEDULE_NOT_IN_DESTINATION**

  A schedule used by a mirror in the pool is not associated with the target domain.
- **DOMAIN_USED_TARGET_NOT_IN_DESTINATION**

  A target that is used by mirror in the pool is not associated with the target domain.
- **DOMAIN_MAX_MIRRORS_REACHED**

  The domain exceeds the maximum allowed number of mirrors.
- **DOMAIN_MAX_DMS_REACHED**

  The domain exceeds the maximum allowed number of data migrations.
- **DOMAIN_MAX_CONS_GROUPS_REACHED**

  The domain exceeds the maximum allowed number of consistency groups.
- **DOMAIN_MAX_VOLUMES_REACHED**

  The domain exceeds the maximum allowed number of volumes.
- **OPERATION_DENIED_OBJECT_MANAGED**

  This is a managed object. Only the managing software and xiv_maintenance / xiv_development may perform this operation on this object.
- **SOURCE_AND_DESTINATION_DOMAINS_MUST_BE_DIFFERENT**

  The source and destination domains must be different.
- **DOMAIN_CONTAINS_OLVM_VOLUME**

  The domain contains a volume in the OLVM process.
- **NO_SPACE**

The system does not have enough free space for the requested storage pool size.

## Setting the domain attribute

Use the **domain_manage** command to set or clear the Managed attribute of a domain.

```
domain_manage domain=DomainName managed=<yes|no>
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **domain** | Object name | The domain name. | Y |
| **managed** | Boolean | The Managed attribute. If set to Yes, the domain will be marked as managed. | Y |

**Example:**

```
domain_mmanage domain=d1 managed=yes
```

**Output:**

```
Command completed successfully.
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Disallowed |
| Storage integration administrator | Allowed |
| Application administrator | Disallowed |
| Security administrator | Disallowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

### Return codes

- **DOMAIN_DOESNT_EXIST**

  The domain does not exist.
- **DOMAIN_HAS_POOL**

  One or more pools are still in the domain.
- **DOMAIN_HAS_SCHEDULE**

  One or more schedules are associated with this domain.
- **DOMAIN_HAS_DEST**

  One or more destinations are associated with this domain.
- **DOMAIN_HAS_DESTGROUP**

  One or more destination groups are associated with this domain.
- **DOMAIN_HAS_USER_GROUP**

  One or more user groups are associated with this domain.

# Setting domain related policies

Use the **domain_policy_set** command to set domain-related policies.

```
domain_policy_set name=Name value=ParamValue
```

## Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **name** | String | Name of the parameter to set. | Y |
| **value** | String | Value of the parameter. | Y |

This command is used for setting domain related policies.

- **name=access** defines whether non-domain-specific users can access domain-specific resources (*value=open*) or not (*value=closed*).
- **name=host_management** defines whether domain administrators can create their own hosts (*value=extended*), or are restricted to hosts assigned to their domains by NDSO administrators (*value=basic*).

**Example:**

```
domain_policy_set name=access value=open
```

**Output:**

```
Command completed successfully.
```

## Access control

| User Category | Permission |
|---|---|
| Storage administrator | Disallowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Allowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

## Return codes

- **UNRECOGNIZED_CONFIG_PARAMETER**

  Unrecognized configuration parameter: '*name*'.

  **Troubleshooting:** Use a valid configuration parameter in the command syntax. For the list of valid configuration parameters, see the CLI Reference Guide.

# Displaying domain-related policies

Use the **domain_policy_get** command to display domain-related policies.

```
domain_policy_get [ name=Name ]
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **name** | String | Name of the parameter to get. | N | All parameters. |

- **name=access** defines whether non-domain-specific users can access domain-specific resources (*value=open*) or not (*value=closed*).
- **name=host_management** defines whether domain administrators can create their own hosts (*value=extended*), or are restricted to hosts assigned to their domains by NDSO administrators (*value=basic*).

| Field ID | Field output | Default position |
|---|---|---|
| **name** | Name | 1 |
| **value** | Value | 2 |

**Example:**

```
domain_policy_get name=access
```

**Output:**

```
Name     Value
-------- -------
access   OPEN
```

## Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Allowed |
| Application administrator | Allowed |
| Security administrator | Allowed |
| Read-only users | Allowed |
| Technicians | Allowed |

## Return codes

- **UNRECOGNIZED_CONFIG_PARAMETER**

  Unrecognized configuration parameter: '*name*'.

  **Troubleshooting:** Use a valid configuration parameter in the command syntax. For the list of valid configuration parameters, see the CLI Reference Guide.

- **CONF_SERVER_UNREACHABLE**

  The configuration server is unreachable.

# Specifying a user associated with IBM Hyper-Scale Manager

Use the **designate_msm_user_set** command to specify the name of the user that is associated with the IBM Hyper-Scale Manager.

```
designate_msm_user_set name=UserName
```

## Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **name** | Object name | The designated user. | Y |

This command specifies which XIV user is defined in the IBM Hyper-Scale Manager Server in the activation step. This can be either a local or LDAP user, depending on whether LDAP authentication is used.

**Example:**

```
designate_msm_user_set name=xiv_msms
```

**Output:**

```
Command completed successfully.
```

## Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Allowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

## Return codes

- **USER_NAME_DOES_NOT_EXIST**

  The user name does not exist.
- **USER_IS_ONLY_DOMAIN_ADMIN**

  The user is associated with one or more domains, and cannot view the entire system.

# Retrieving the user associated with the IBM Hyper-Scale Manager

Use the **designate_msm_user_get** command to retrieve the name of the user associated with the IBM Hyper-Scale Manager.

```
designate_msm_user_get
```

**Example:**

```
designate_msm_user_get
```

**Output:**

```
xiv_msms
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Allowed |
| Application administrator | Allowed |
| Security administrator | Allowed |
| Read-only users | Allowed |
| Technicians | Allowed |

### Return codes

- **NO_DESIGNATED_MSM_USER**

  There is no designated IBM Hyper-Scale user.

## Setting the application administrator's scope of commands

Use the **appadmin_capabilities_set** command to define whether an Application Administrator is authorized to perform the basic or advanced set of commands.

```
appadmin_capabilities_set value=<basic|advanced>
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **value** | Enumeration | The set of commands that an Application Administrator is authorized to perform. | Y |

These are the additional commands available when an application administrator is authorized to perform advanced commands:

- **mirror_statistics_get**
- **mirror_activate**
- **mirror_deactivate**
- **mirror_change_role**
- **mirror_switch_roles**

**Example:**

```
appadmin_capabilities_set value=basic
```

**Output:**

```
Command completed successfully.
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Disallowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Allowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

## Getting the application administrator's scope of commands

Use the **appadmin_capabilities_get** command to display the state of the Application Administrator's capabilities.

```
appadmin_capabilities_get
```

**Example:**

```
appadmin_capabilities_get
```

**Output:**

```
BASIC
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Allowed |
| Application administrator | Allowed |
| Security administrator | Allowed |
| Read-only users | Allowed |
| Technicians | Allowed |

# 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 `type=target`, the *non-logged-in* option can only be used to debug targets that rejected the storage system login. This command returns an error for an attempt to list targets from a target-only port, or to list initiators from an initiator-only port. Each output line contains the following information:

- Component ID (of the module)
- Storage system port number (within the module)
- WWPN
- Port ID (can be correlated with the switch database)
- Role: Initiator, Target, Dual
- Initiator/target (is the same for all lines of the same command)
- Login status (Yes/No)

| Field ID | Field output | Default position |
|---|---|---|
| **component_id** | Component ID | 1 |
| **wwpn** | WWPN | 2 |
| **port_id** | Port ID | 3 |
| **role** | Role | 4 |

© Copyright IBM Corp. 2016, 2022

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Allowed |
| Security administrator | Disallowed |
| Read-only users | Allowed |
| Technicians | Allowed |

## Changing FC port configuration

Use the **fc_port_config** command to configure FC ports.

```
fc_port_config fc_port=ComponentId [ enabled=<yes|no> ] [ role=<target|initiator> ] [ rate=<2|4|8|
16|auto> ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **fc_port** | N/A | Port identifier. | Y | N/A |
| **enabled** | Boolean | Allows you to enable or disable the port. | N | yes |
| **role** | Enumeration | Port role: target, initiator or both. | N | Leaves the role unchanged. |
| **rate** | Enumeration | Line rate or auto for auto-negotiated rate. | N | Leaves the rate unchanged. |

**Example:**

```
fc_port_config fc_port=1:FC_Port:1:1 enabled=yes role=Target rate=auto
```

**Output:**

```
Command completed successfully
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Disallowed |
| Read-only users | Disallowed |
| Technicians | Allowed |

### Return codes

- **COMPONENT_IS_NOT_AN_FC_PORT**

  An FC port must be specified for the component.

- **FC_PORT_DOES_NOT_EXIST**

  The specified FC port does not exist.
- **FC_PORT_TEST_IN_PROGRESS**

  The FC port test is already in progress

## Listing FC ports

Use the **fc_port_list** command to list the status and configuration of the system's FC ports.

```
fc_port_list [ module=ModuleNumber  | fcport=ComponentId ]
```

### Parameters

| Name | Description | Mandatory | Default |
|---|---|---|---|
| **module** | Limits the listing to a specific module. | N | All ports in all modules. |
| **fcport** | Lists only a specific port. | N | All ports in all modules. |

This command lists all or some FC ports on the system. When no parameters are specified, all ports are listed. If a module is specified without a port, all ports on that module are listed. If a port is specified, a single port is listed.

The following information is provided for each port:

- Component ID of the module Port number (internal to module) 1-N
- WWPN
- Port ID
- Role (Initiator, Target, Dual)
- User-enabled (Yes/No)
- Maximum support rate: 2GB, 4GB, 8GB; constant - function of the HBA's capability
- Configured rate: 2GB, 4GB, 8GB, auto-negotiation; cannot be greater than the maximum supported rate
- Current active rate: 2GB, 4GB, 8GB; equal to the configured rate, unless the configured rate is auto-negotiation
- Port state: Online, Offline, Loopback, Link Down (physical connection is on, but no logical connection exists)
- Error counts
- Link type: Fabric Direct Attach, Private Loop, Point-to-Point, Public Loop, Unknown

**Example:**

```
fc_port_list
```

**Output:**

```
Component ID     Status   Currently Functioning   WWPN               Port ID    Role
---------------- -------- ---------------------- ------------------ ---------- --------
1:FC_Port:12:1   OK       yes                    5001738035C601C0   FFFFFFFF   Target
1:FC_Port:12:2   OK       yes                    5001738035C601C1   FFFFFFFF   Target
1:FC_Port:12:3   OK       yes                    5001738035C601C2   FFFFFFFF   Target
1:FC_Port:12:4   OK       yes                    5001738035C601C3   00EF009A   Target
1:FC_Port:13:1   OK       yes                    5001738035C601D0   FFFFFFFF   Target
1:FC_Port:13:2   OK       yes                    5001738035C601D1   FFFFFFFF   Target
1:FC_Port:13:3   OK       yes                    5001738035C601D2   FFFFFFFF   Target
1:FC_Port:13:4   OK       yes                    5001738035C601D3   FFFFFFFF   Target
1:FC_Port:8:1    OK       yes                    5001738035C60180   FFFFFFFF   Target
1:FC_Port:8:2    OK       yes                    5001738035C60181   FFFFFFFF   Target
1:FC_Port:8:3    OK       yes                    5001738035C60182   FFFFFFFF   Target
1:FC_Port:8:4    OK       yes                    5001738035C60183   00163AC0   Target

Cont.:

User Enabled   Current Rate (GBaud)   Port State     Link Type
-------------- ---------------------- -------------- ----------------------
yes            Auto                   Link Problem   None
yes            Auto                   Link Problem   None
yes            Auto                   Link Problem   None
yes            8                      Online         Fabric Direct Attach
yes            Auto                   Link Problem   None
yes            Auto                   Link Problem   None
yes            Auto                   Link Problem   None
yes            Auto                   Link Problem   None
yes            Auto                   Link Problem   None
yes            Auto                   Link Problem   None
yes            Auto                   Link Problem   None
yes            16                     Online         Fabric Direct Attach

Cont.:

Error Count   Active Firmware
------------- -----------------
0             8.3.40
0             8.3.40
0             8.3.40
0             8:3:40
0             8.3.40
0             8.3.40
0             8.3.40
0             8.3.40
0             8.3.40
0             8.3.40
0             8.3.40
0             8:3:40
```

| Field ID | Field output | Default position |
|---|---|---|
| **component_id** | Component ID | 1 |
| **status** | Status | 2 |
| **currently_functioning** | Currently Functioning | 3 |
| **port_num** | Port Number | N/A |
| **wwpn** | WWPN | 4 |
| **port_id** | Port ID | 5 |
| **role** | Role | 6 |
| **user_enabled** | User Enabled | 7 |
| **max_supported_rate** | Maximum Supported Rate (GBaud) | N/A |
| **configured_rate** | Configured Rate (GBaud) | N/A |
| **current_rate** | Current Rate (GBaud) | 8 |
| **port_state** | Port State | 9 |
| **link_type** | Link Type | 10 |
| **error_count** | Error Count | 11 |

| Field ID | Field output | Default position |
|---|---|---|
| **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 on the system. When no parameters are specified, all tests are listed. If a port is specified, a single test is listed.

**Example:**

```
fc_port_test_list
```

**Output:**

```
Component ID     Status       Start Time   Duration (Sec)   Transfer Length   CRC
---------------- ------------ ------------ ---------------- ----------------- -----
1:FC_Port:14:1   not active   NA           0                0                 0
1:FC_Port:14:2   not active   NA           0                0                 0
1:FC_Port:14:3   not active   NA           0                0                 0
1:FC_Port:14:4   not active   NA           0                0                 0
1:FC_Port:3:1    not active   NA           0                0                 0
1:FC_Port:3:2    not active   NA           0                0                 0
1:FC_Port:3:3    not active   NA           0                0                 0
1:FC_Port:3:4    not active   NA           0                0                 0
1:FC_Port:5:1    not active   NA           0                0                 0
1:FC_Port:5:2    not active   NA           0                0                 0
1:FC_Port:5:3    not active   NA           0                0                 0
1:FC_Port:5:4    not active   NA           0                0                 0
1:FC_Port:6:1    not active   NA           0                0                 0
1:FC_Port:6:2    not active   NA           0                0                 0
1:FC_Port:6:3    not active   NA           0                0                 0
1:FC_Port:6:4    not active   NA           0                0                 0

Cont.:

Frame Length Error   Disparity   Loopback Mode   Data Pattern   Data Size   Test Count
-------------------- ----------- --------------- -------------- ----------- ------------
0                    0           internal        0000           0           0
0                    0           internal        0000           0           0
0                    0           internal        0000           0           0
0                    0           internal        0000           0           0
0                    0           internal        0000           0           0
0                    0           internal        0000           0           0
0                    0           internal        0000           0           0
0                    0           internal        0000           0           0
0                    0           internal        0000           0           0
0                    0           internal        0000           0           0
0                    0           internal        0000           0           0
0                    0           internal        0000           0           0
0                    0           internal        0000           0           0
0                    0           internal        0000           0           0
0                    0           internal        0000           0           0
0                    0           internal        0000           0           0

Cont.:

Abort On Error
----------------
no
no
no
no
no
no
no
no
no
no
no
no
no
no
no
no
no
no
```

| Field ID | Field output | Default position |
|---|---|---|
| **component_id** | Component ID | 1 |
| **status** | Status | 2 |
| **start_time** | Start Time | 3 |

| Field ID | Field output | Default position |
|---|---|---|
| **duration** | Duration (Sec) | 4 |
| **transfer_length** | Transfer Length | 5 |
| **crc** | CRC | 6 |
| **frame_length_error** | Frame Length Error | 7 |
| **disparity** | Disparity | 8 |
| **loopback_mode** | Loopback Mode | 9 |
| **pattern** | Data Pattern | 10 |
| **data_size** | Data Size | 11 |
| **test_count** | Test Count | 12 |
| **increment** | Test Increment | N/A |
| **abort_on_error** | Abort On Error | 13 |

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Allowed |
| Application administrator | Allowed |
| Security administrator | Disallowed |
| Read-only users | Allowed |
| Technicians | Allowed |

## Starting FC Port Test

Use **fc_port_test_start** to run a test of the FC port.

```
fc_port_test_start fc_port=ComponentId [ loopback=<internal|external> ] [ pattern=PATTERN ]
 [ data_size=DATA_SIZE ] [ frames=FRAMES ] [ increment=INCREMENT ]
 [ abort_on_error=<no|yes> ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **fc_port** | N/A | Port identifier. | Y | N/A |
| **loopback** | Enumeration | Loopback mode to use. | N | internal |
| **pattern** | String | Data pattern. | N | 1414 |
| **data_size** | Positive integer | Data size. | N | 2048 |
| **frames** | Positive integer | Number of frames. | N | 10000 |
| **increment** | Positive integer | Test increment. | N | 1 |
| **abort_on_error** | Boolean | Abort the test in case of an error. | N | no |

**Example:**

```
fc_port_test_start fc_port=1:FC_Port:1:1
```

**Output:**

```
Command completed successfully
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Disallowed |
| Read-only users | Disallowed |
| Technicians | Allowed |

### Return codes

- **COMPONENT_IS_NOT_AN_FC_PORT**

  An FC port must be specified for the component.
- **FC_PORT_DOES_NOT_EXIST**

  The specified FC port does not exist.
- **COMPONENT_DOES_NOT_EXIST**

  The component does not exist.
- **TEST_NOT_ALLOWED_IN_CURRENT_STATUS**

  This component cannot be tested in its current status.
- **FC_PORT_TEST_IN_PROGRESS**

  The FC port test is already in progress
- **INVALID_FC_PORT_TEST_DATA_PATTERN**

  The data pattern for the FC port test is invalid.

## Aborting FC Port Test

Use **fc_port_test_abort** to abort a currently running FC port test.

```
fc_port_test_abort fc_port=ComponentId
```

### Parameters

| Name | Description | Mandatory |
|---|---|---|
| **fc_port** | Port identifier. | Y |

**Example:**

```
fc_port_test_abort fc_port=1:FC_Port:1:1
```

**Output:**

```
Command completed successfully
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Disallowed |
| Read-only users | Disallowed |
| Technicians | Allowed |

### Return codes

- **COMPONENT_IS_NOT_AN_FC_PORT**

  An FC port must be specified for the component.
- **FC_PORT_DOES_NOT_EXIST**

  The specified FC port does not exist.
- **COMPONENT_DOES_NOT_EXIST**

  The component does not exist.
- **NO_FC_PORT_TEST_IN_PROGRESS**

  The FC port test is currently not running.

## Listing connectivity to hosts

Use the **host_connectivity_list** command to list Fibre Channel and iSCSI-level connectivity.

```
host_connectivity_list [ host=HostName | fc_host_port=WWPN ] [ module=ModuleNumber |
fcport=ComponentId ] [ domain=DomainName ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **host** | Object name | Limit list to the ports of a specific host. | N | All hosts. |
| **fc_host_port** | N/A | Limit list to this specific port. | N | All ports |
| **module** | N/A | Limits output only to the enabled connectivity to this module. | N | All modules |
| **fcport** | N/A | Limits output to a specific storage system's port. | N | All ports |
| **domain** | Object name | The domain name. | N | All Domains |

This command shows the connectivity status between a storage system port and a defined host. The output can be limited to a specific port, module or storage system port. Hosts can attach to the FC and iSCSI either directly (point-to-point), via a Fibre Channel fabric, or via a Gigabit Ethernet switch. Connectivity refers to both physical connectivity and SCSI login. Each output line contains the following information:

- Host (name)
- Host port (WWPN)

- Module ID, preceded by the rack ID
- Port number (within the module)

**Example:**

```
host_connectivity_list host=demo_host_fc0 fc_host_port=1:FC_Port:5:1
```

**Output:**

```
Host          Host Port        Module      Local FC port   Local iSCSI port   Type
-----         ----------       ----------- --------------- ----------------   -----
demo_host_fc0 100000062B151A98 1:Module:5  1:FC_Port:5:1                      FC
```

| Field ID | Field output | Default position |
|---|---|---|
| **host** | Host | 1 |
| **host_port** | Host Port | 2 |
| **module** | Module | 3 |
| **local_fc_port** | Local FC port | 4 |
| **local_iscsi_port** | Local iSCSI port | 5 |
| **type** | Type | 6 |
| **vlan_id** | VLAN ID | 7 |
| **local_ipinterface** | Local IP interface | 8 |
| **local_ethernet_port** | Local Ethenet port | 9 |

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Allowed |
| Application administrator | Allowed |
| Security administrator | Disallowed |
| Read-only users | Allowed |
| Technicians | Allowed |

## 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
```

| Field ID | Field output | Default position |
|---|---|---|
| **module_name** | Module | 1 |
| **module_1_link_status** | module-1 | 2 |
| **module_2_link_status** | module-2 | 3 |
| **module_3_link_status** | module-3 | 4 |
| **canister_1_link_status** | Flash-Canister-1 | 5 |
| **canister_2_link_status** | Flash-Canister-2 | 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 |

# 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   Cluster IP
--------------------- -------- ----------------------- --------------------- -------------
1:Flash_Enclosure:1   OK       yes                     OK                    14.10.204.3

Cont.:

Redundancy State   FW level         Has Spare   Array Rebuild Percentage
------------------ ---------------- ----------- -------------------------
online             1.5.0.0-436.13   yes         None

Cont.:

Encryption State   Machine Model
------------------ ---------------
Ready              AE2
```

| Field ID | Field output | Default position |
|---|---|---|
| **component_id** | Component ID | 1 |
| **status** | Status | 2 |
| **currently_functioning** | Currently Functioning | 3 |
| **control_path_status** | Control Path Status | 4 |
| **cluster_ip** | Cluster IP | 5 |
| **redundancy_state** | Redundancy State | 6 |
| **fw_level** | FW level | 7 |
| **has_spare** | Has Spare | 8 |

© Copyright IBM Corp. 2016, 2022

| Field ID | Field output | Default position |
|---|---|---|
| **array_rebuild_percentage** | Array Rebuild Percentage | 9 |
| **machine_model** | Machine Model | 11 |
| **array_status** | Array Status | N/A |
| **fru_part_number** | FRU Part Number | N/A |
| **fru_identity** | FRU Identity | N/A |
| **temperature_state** | Temperature State | N/A |
| **required_service** | Requires Service | N/A |
| **service_reason** | Service Reason | N/A |
| **enabled** | Enabled | N/A |
| **cluster_id** | Cluster ID | N/A |
| **serial_number** | Serial Number | N/A |
| **encrypted** | Encrypted | N/A |
| **key_needed** | Key Needed | N/A |
| **encryption_state** | Encryption State | 10 |
| **base_guid** | Base GUID | N/A |
| **charging** | Charging | N/A |
| **flash_status** | Flash Status | N/A |
| **fw_upgrade_status** | FW Upgrade Status | N/A |
| **fw_upgrade_progress** | FW Upgrade Progress | N/A |
| **target_fw_version** | Target FW. Version | N/A |
| **fw_file_name** | FW File Name | N/A |
| **utility_file_name** | Utility File Name | N/A |
| **cr_key_last_modified_time** | CR Key Last Modified Time | N/A |

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Allowed |
| Security administrator | Allowed |
| Read-only users | Allowed |
| Technicians | Allowed |

## Listing a Flash canister status

Use the **flash_canister_list** command to list special Flash canister statuses.

```
flash_canister_list [ canister=ComponentId | flash_enclosure=ComponentId ]
```

## Parameters

| Name | Description | Mandatory | Default |
|---|---|---|---|
| **canister** | Flash Canister for which special statuses are to be listed. | N | All Flash canisters. |
| **flash_enclosure** | Flash Enclosure for which special statuses are to be listed. | N | All Flash enclosures. |

This command lists the statuses of the Flash canisters, including:

- Component generic status
- Canister ID
- Node ID
- Node name

**Example:**

```
flash_canister_list
```

**Output:**

```
Component ID           Status   Currently Functioning   Service IP      Raid Status
---------------------- -------- ----------------------- --------------- -------------
1:Flash_Canister:4:1   OK       yes                     14.10.204.77    OK
1:Flash_Canister:4:2   OK       yes                     14.10.204.110   OK

Cont.:

Control Path Status   Serial Connected   MgmtNode
--------------------- ------------------ ----------
OK                    1:Module:12        no
OK                    1:Module:13        yes
```

| Field ID | Field output | Default position |
|---|---|---|
| **component_id** | Component ID | 1 |
| **status** | Status | 2 |
| **currently_functioning** | Currently Functioning | 3 |
| **service_ip** | Service IP | 4 |
| **raid_status** | Raid Status | 5 |
| **control_path_status** | Control Path Status | 6 |
| **serial_connected** | Serial Connected | 7 |
| **active** | MgmtNode | 8 |
| **fru_part_number** | FRU Part Number | N/A |
| **fru_identity** | FRU Identity | N/A |
| **temperature_state** | Temperature State | N/A |
| **fw_level** | FW Level | N/A |
| **mac_addresses** | MAC Addresses | N/A |
| **required_service** | Requires Service | N/A |
| **service_reason** | Service Reason | N/A |
| **canister_degraded** | Canister Degraded | N/A |
| **canister_missing** | Canister Missing | N/A |

| Field ID | Field output | Default position |
|---|---|---|
| **status_led** | Status LED | N/A |
| **check_log_led** | Check Log LED | N/A |
| **identify_led** | Identify LED | N/A |
| **controller_fault_led** | Controller Fault LED | N/A |
| **fault_led** | Fault LED | N/A |
| **dump_led** | Dump LED | N/A |
| **canister_mode** | Canister Mode | N/A |
| **service_mode** | Service Mode | N/A |
| **miswired** | Miswired | N/A |

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Allowed |
| Security administrator | Allowed |
| Read-only users | Allowed |
| Technicians | Allowed |

## Listing a Flash card status

Use the **flash_card_list** command to list special Flash card statuses.

```
flash_card_list [ flash_card=ComponentId | flash_enclosure=ComponentId ]
```

### Parameters

| Name | Description | Mandatory | Default |
|---|---|---|---|
| **flash_card** | Flash card for which special statuses are to be listed. | N | All Flash cards. |
| **flash_enclosure** | Flash enclosure for which special statuses are to be listed. | N | All Flash enclosures. |

This command lists the statuses of a Flash card, including:

- Component generic status
- Slot ID
- Capacity
- Health State
- Flash type

**Example:**

```
flash_card_list flash_card=1:Flash_Card:4:5
```

**Output:**

```
Component ID       Status   Currently Functioning   Slot ID   Capacity(GB)   Health State
-----------------  -------- ----------------------  --------- -------------  --------------
1:Flash_Card:4:5   OK       yes                     5         5717           good

Cont.:

Usage    Missing
-------- ---------
member   no
```

| Field ID | Field output | Default position |
|---|---|---|
| **component_id** | Component ID | 1 |
| **status** | Status | 2 |
| **currently_functioning** | Currently Functioning | 3 |
| **slot_id** | Slot ID | 4 |
| **capacity** | Capacity(GB) | 5 |
| **health_state** | Health State | 6 |
| **drive_use** | Usage | 7 |
| **missing** | Missing | 8 |
| **fru_part_number** | FRU Part Number | N/A |
| **fru_identity** | FRU Identity | N/A |
| **temperature_state** | Temperature State | N/A |
| **fw_level** | FW Level | N/A |
| **required_service** | Requires Service | N/A |
| **service_reason** | Service Reason | N/A |
| **flash_type** | Type | N/A |
| **fault_led** | Fault LED | N/A |

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Allowed |
| Security administrator | Allowed |
| Read-only users | Allowed |
| Technicians | Allowed |

## Retrieving the list of Flash fans

Use the **flash_fan_list** command to retrieve the list of Flash fans.

```
flash_fan_list  [ flash_fan=ComponentId | flash_enclosure=ComponentId ]
```

### Parameters

| Name | Description | Mandatory | Default |
|---|---|---|---|
| **flash_fan** | Flash fan component ID | N | all |

| Name | Description | Mandatory | Default |
|---|---|---|---|
| **flash_enclosure** | Flash enclosure component ID | N | all |

**Example:**

```
flash_fan_list
```

**Output:**

```
Component ID      Status
----------------- --------
1:Flash_Fan:2:1   OK
1:Flash_Fan:2:2   OK
1:Flash_Fan:2:3   OK
1:Flash_Fan:2:4   OK
```

| Field ID | Field output | Default position |
|---|---|---|
| **component_id** | Component ID | 1 |
| **status** | Status | 2 |
| **currently_functioning** | Currently Functioning | 3 |
| **fru_part_number** | FRU Part Number | N/A |
| **fru_identity** | FRU Identity | N/A |
| **temperature_state** | Temperature State | N/A |
| **fw_level** | FW Level | N/A |
| **required_service** | Requires Service | N/A |
| **service_reason** | Service Reason | N/A |

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Allowed |
| Security administrator | Disallowed |
| Read-only users | Allowed |
| Technicians | Allowed |

## Retrieving the list of Flash IB adapters

Use the **flash_ib_adapter_list** command to retrieve the list of Flash IB adapters.

```
flash_ib_adapter_list  [ flash_ib_adapter=ComponentId | flash_enclosure=ComponentId ]
```

### Parameters

| Name | Description | Mandatory | Default |
|---|---|---|---|
| **flash_ib_adapter** | Flash IB adapter component ID | N | all |
| **flash_enclosure** | Flash Enclosure component ID | N | all |

**Example:**

```
flash_ib_adapter_list
```

**Output:**

```
Component ID             Status
------------------------ --------
1:Flash_IB_Adapter:2:1   OK
1:Flash_IB_Adapter:2:2   OK
1:Flash_IB_Adapter:2:3   OK
1:Flash_IB_Adapter:2:4   OK
```

| Field ID | Field output | Default position |
|---|---|---|
| **component_id** | Component ID | 1 |
| **status** | Status | 2 |
| **currently_functioning** | Currently Functioning | 3 |
| **canister_id** | Canister_name | 4 |
| **fru_part_number** | FRU Part Number | N/A |
| **fru_identity** | FRU Identity | N/A |
| **temperature_state** | Temperature State | N/A |
| **required_service** | Requires Service | N/A |
| **service_reason** | Service Reason | N/A |
| **fw_level** | FW Level | N/A |
| **port1_id** | Port 1 Component ID | N/A |
| **port1_guid** | Port 1 GUID | N/A |
| **port2_id** | Port 2 Component ID | N/A |
| **port2_guid** | Port 2 GUID | N/A |

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Allowed |
| Security administrator | Disallowed |
| Read-only users | Allowed |
| Technicians | Allowed |

## Retrieving the Flash control connectivity list

Use the **flash_control_connectivity_list** command to retrieve the Flash control connectivity list.

```
flash_control_connectivity_list  [ canister=ComponentId | flash_enclosure=ComponentId ]
```

## Parameters

| Name | Description | Mandatory | Default |
|---|---|---|---|
| **canister** | Canister component ID | N | all |
| **flash_enclosure** | Flash enclosure Component ID | N | all |

This command is used to list the Ethernet connectivity to flash canisters in the system. On FlashSystem A9000R systems, all flash canisters are connected via grid controllers (modules) 1 through 4. On FlashSystem A9000 systems, each flash canister is connected directly to a specific grid controller.

The command output lists each flash canister and its connectivity to each of the grid controllers. For each flash canister, the path to each grid controller includes the flash canister IP address used by the grid controller and the connectivity status.

**Example:**

```
flash_control_connectivity_list
```

**Output:**

```
This command's output is dependent on the FlashSystem type

Output for FlashSystem A9000R:

Component ID           Path1        PathAddr1      P1Status  Path2        PathAddr2
---------------------- ------------ -------------- ---------- ------------ --------------
1:Flash_Canister:1:1   1:Module:1   14.10.204.1    OK        1:Module:2   14.10.204.33
1:Flash_Canister:1:2   1:Module:1   14.10.204.2    OK        1:Module:2   14.10.204.34
1:Flash_Canister:2:1   1:Module:1   14.10.204.4    OK        1:Module:2   14.10.204.37
1:Flash_Canister:2:2   1:Module:1   14.10.204.5    OK        1:Module:2   14.10.204.38

P2Status   Path3        PathAddr3      P3Status   Path4        PathAddr4       P4Status
---------- ------------ -------------- ---------- ------------ --------------- ----------
OK         1:Module:3   14.10.204.65   OK         1:Module:4   14.10.204.97    OK
OK         1:Module:3   14.10.204.66   OK         1:Module:4   14.10.204.98    OK
OK         1:Module:3   14.10.204.69   OK         1:Module:4   14.10.204.101   OK
OK         1:Module:3   14.10.204.70   OK         1:Module:4   14.10.204.102   OK


Output for FlashSystem A9000:

Component ID           Path1        PathAddr1      P1Status   Path2        PathAddr2
---------------------- ------------ -------------- ---------- ------------ --------------
1:Flash_Canister:1:1                                N.A.       1:Module:2   14.10.204.1
1:Flash_Canister:1:2   1:Module:1   14.10.204.34   OK

P2Status   Path3        PathAddr3      P3Status   Path4        PathAddr4       P4Status
---------- ------------ -------------- ---------- ------------ --------------- ----------
OK                                     N.A.                                   N.A.
N.A.                                   N.A.                                   N.A.
```

| Field ID | Field output | Default position |
|---|---|---|
| **component_id** | Component ID | 1 |
| **virtual_ips.0.path** | Path1 | 2 |
| **virtual_ips.0.pathAddr** | PathAddr1 | 3 |
| **virtual_ips.0.status** | P1Status | 4 |
| **virtual_ips.0.state** | P1State | N/A |
| **virtual_ips.1.path** | Path2 | 5 |
| **virtual_ips.1.pathAddr** | PathAddr2 | 6 |
| **virtual_ips.1.status** | P2Status | 7 |
| **virtual_ips.1.state** | P2State | N/A |

| Field ID | Field output | Default position |
|---|---|---|
| **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 |

| Field ID | Field output | Default position |
|---|---|---|
| **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 |

| Field ID | Field output | Default position |
|---|---|---|
| **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 |

| Field ID | Field output | Default position |
|---|---|---|
| **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 |

| Field ID | Field output | Default position |
|---|---|---|
| **temperature_state** | Temperature State | N/A |
| **fw_level** | FW Level | N/A |
| **required_service** | Requires Service | N/A |
| **service_reason** | Service Reason | N/A |
| **fru_part_number** | FRU Part Number | N/A |
| **fru_identity** | FRU Identity | N/A |
| **fault_led** | Fault LED | N/A |

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Allowed |
| Security administrator | Disallowed |
| Read-only users | Allowed |
| Technicians | Allowed |

## Retrieving the serial number of a flash enclosure 1S

Use the **flash_enclosure_1s_get** command to retrieve the serial number of a flash enclosure 1S.

```
flash_enclosure_1s_get flash_enclosure=ComponentId
```

### Parameters

| Name | Description | Mandatory |
|---|---|---|
| **flash_enclosure** | Flash enclosure component ID | Y |

**Example:**

```
flash_enclosure_1s_get
```

| Field ID | Field output | Default position |
|---|---|---|
| **full_serial** | Serial | 1 |
| **mtm** | Mtm | N/A |
| **serial** | Serial | N/A |

**Example:**

```
flash_enclosure_1s_get
```

**Output:**

```
Serial
----------------
98354151234567
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Allowed |
| Security administrator | Disallowed |
| Read-only users | Allowed |
| Technicians | Allowed |

### Return codes

- **COMMAND_IS_NOT_VALID_DURING_FLASH_FW_UPDATE**

  The requested command cannot be invoked while a FlashSystem update is running.
- **FLASH_ENCLOSURE_DOES_NOT_EXIST**

  Flash enclosure does not exist in the system

## Suspending the automatic Flash BBU calibration

Use the **flash_system_bbu_calibration_set** command to suspend and resume the automatic Flash BBU calibration.

```
flash_system_bbu_calibration_set suspend=<yes|no> [ suspend_days=days ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **suspend** | Boolean | Suspend calibration | Y | N/A |
| **suspend_days** | Positive integer | Days to suspend | N | 14 |

**Example:**

```
flash_system_bbu_calibration_set suspend=yes suspend_days=14
```

**Output:**

```
Command completed successfully.
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Allowed |
| Security administrator | Disallowed |
| Read-only users | Disallowed |
| Technicians | Allowed |

# Chapter 21. Hardware maintenance commands

This section describes the command-line interface (CLI) for maintaining hardware components.

## Monitoring the rebuilding or redistribution processes

Use the **monitor_redist** command to monitor the status of the rebuilding or redistribution process.

```
monitor_redist
```

This command outputs the current redistribution process running on a module. The command does not yield information about a Flash enclosure.

The command may inform you that no such process exists. If such a process exists, the following information is shown:

- Type (adding new capacity, replacing failed component, phase-out, rebuild after failure)
- Initial capacity to copy
- Time started
- Capacity remaining to copy
- Time elapsed
- Percent completed
- Estimated time to completion

| Field ID | Field output | Default position |
|---|---|---|
| **type** | Type | 1 |
| **partitions_total** | Initial Partitions | 2 |
| **partitions_left** | Partitions Remaining | 3 |
| **percent_done** | % Done | 4 |
| **time_started** | Time Started | 5 |
| **estimated_time_to_finish** | Estimated Time to Finish | 6 |
| **time_elapsed** | Time Elapsed | 7 |

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Allowed |
| Security administrator | Disallowed |
| Read-only users | Allowed |
| Technicians | Allowed |

© Copyright IBM Corp. 2016, 2022

# Displaying the system's average power consumption

Use the **system_average_power_consumption** command to display the system's average power consumption.

```
system_average_power_consumption
```

| Field ID | Field output | Default position |
|---|---|---|
| **value** | Value | 1 |
| **timestamp** | Timestamp | 2 |

**Note:**

The average power consumption value is indicated in Watts.

**Example:**

```
system_average_power_consumption
```

**Output:**

```
Value       Timestamp
----------- -----------
1337W       17-Jan-2020 12:30:27
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Allowed |
| Security administrator | Disallowed |
| Read-only users | Allowed |
| Technicians | Allowed |

# Getting the values for calculating the system's average power consumption

Use the **system_average_power_prepare** command to fetch all the values needed to calculate the system's average power consumption.

```
system_average_power_prepare
```

**Example:**

```
system_average_power_prepare
```

**Output:**

```
Command completed successfully
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Disallowed |
| Read-only users | Disallowed |
| Technicians | Allowed |

### Return codes

- **HOT_UPGRADE_IS_IN_PROGRESS**

  The command is not allowed during hot upgrade.

  **Troubleshooting:** Wait for the hot upgrade to complete and try again.
- **ALREADY_GETTING_FLASH_CANISTER_SNAP**

  Already getting a flash canister snap.
- **FAILED_GETTING_FLASH_CANISTER_SNAP**

  Failed getting a flash canister snap.
- **COMPONENT_DOES_NOT_EXIST**

  The component does not exist.
- **COMMAND_IS_NOT_VALID_DURING_FLASH_FW_UPDATE**

  The requested command cannot be invoked while a FlashSystem update is running.
- **SYSTEM_POWER_PREPARE_ALREADY_IN_PROGRESS**

  There is already a system power prepare command in progress.

## Displaying the system's average temperature

Use the **system_average_temperature** command to display the system's average temperature.

```
system_average_temperature
```

| Field ID | Field output | Default position |
|---|---|---|
| **value** | Value | 1 |
| **timestamp** | Timestamp | 2 |

**Note:**

The average temperature value is indicated in Celsius.

**Example:**

```
system_average_temperature
```

**Output:**

```
Value       Timestamp
----------- --------------------
17C         6-Feb-2020 12:15:22
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Allowed |
| Security administrator | Disallowed |
| Read-only users | Allowed |
| Technicians | Allowed |

## Enabling XIV Support access

Use the **xiv_support_enable** command to enable XIV Support access for a specific period of time limiting access from the specific address.

```
xiv_support_enable [ start=TimeStamp ] < finish=TimeStamp | timeout=Timeout >
from=<IPaddress1[,IPaddress2]
...> comment=Comment
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **start** | N/A | Start time for allowing XIV Support access. | N | Immediately. |
| **finish** | N/A | End time for allowing XIV Support access. | N | N/A |
| **timeout** | N/A | Timeout for allowing XIV Support access in either hh:mm format, or a number of minutes. The timeout cannot exceed 23 hours and 59 minutes. The word `unlimited` denotes unexpired timeout. | N | N/A |
| **from** | N/A | The source address to which XIV Support access is limited. It may be either IPv4 or IPv6 address, or `any`, or `technician` denoting notebook port. | Y | N/A |
| **comment** | String | Reason why XIV Support access is enabled. | Y | N/A |

This command enables XIV Support access for a specific period of time limiting access from the specific address.

**Example:**

```
xiv_support_enable finish=2012-2-3.16:30 from=192.0.2.1 comment="Some reason"
```

**Output:**

```
Command completed successfully.
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Allowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

### Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_ENABLE_XIV_SUPPORT**

  Are you sure you want to enable XIV support?

### Return codes

- **XIV_SUPPORT_WORK_INVALID_TIMEOUT**

  Timeout must be a positive number and define a time greater then the current time
- **XIV_SUPPORT_WORK_INVALID_FINISH**

  The end time must be greater than the start time and the current time.
- **XIV_SUPPORT_WORK_INVALID_FROM**

  From must be a valid IPv4 or IPv6 address.
- **LIST_WITH_MIXED_IPV6_AND_IPV4_NOT_ALLOWED**

  All IP addresses in the list should be of the same type - either IPv4 or IPv6.
- **LIST_WITH_ANY_OPTION_AND_SPECIFIC_IP_ADDRESSES_NOT_ALLOWED**

  All IP addresses in the list should be unicast or Any. Mixing unicast and Any in the same list is not allowed.

## Disabling XIV Support access

Use the **xiv_support_disable** command to disable XIV Support access.

```
xiv_support_disable
```

**Example:**

```
xiv_support_disable
```

**Output:**

```
Command completed successfully.
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |

| User Category | Permission |
|---|---|
| Application administrator | Disallowed |
| Security administrator | Allowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

# Displaying the XIV Support window

Use the **xiv_support_show** command to display the XIV Support window.

```
xiv_support_show
```

The following information is listed:

- From (IPv4 or IPv6 addresses, or "any address", or "technician port")
- Start (timestamp or "unlimited")
- Finish (timestamp or "unlimited")
- Comment

**Example:**

```
xiv_support_show
```

**Output:**

```
Enabled   Start                Finish               Comment
--------- -------------------- -------------------- -----------
yes       2012-03-28 12:55:21  2012-03-30 00:00:00  some work

cont:
From 0   From 1   From 2   From 3   From 4
-------- -------- -------- -------- --------
192.0.2.1
```

| Field ID | Field output | Default position |
|---|---|---|
| **enabled** | Enabled | 1 |
| **start** | Start | 2 |
| **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 |

| User Category | Permission |
|---|---|
| Application administrator | Allowed |
| Security administrator | Allowed |
| Read-only users | Allowed |
| Technicians | Allowed |

## Listing system components

Use the **component_list** command to list system components and their status.

```
component_list [ component=ComponentId ] [ filter=<ALL|FAILED|NOTOK> ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **component** | N/A | Lists only this component. | N | All components. |
| **filter** | Enumeration | Filters the list to show only failed or only non-OK components. | N | ALL |

Lists system components. The list can be filtered to show only a specific component, all failed components or all components in a non-OK state.

For status and configuration of specific component types, refer to the **_list** commands for specific components, such as: **module_list** or **switch_list**.

**Example:**

```
component_list
```

**Output:**

```
Component ID                 Status         Currently Functioning
---------------------------- -------------- ----------------------
1:BBU:12:1                   Failed         no
1:Boot_Media:12:1            OK             yes
1:Boot_Media:12:2            OK             yes
1:CNA:12:1                   OK             yes
1:CNA:13:1                   OK             yes
1:CNA:8:1                    OK             yes
1:CPU:12:1                   OK             yes
1:CPU:12:2                   OK             yes
1:DIMM:12:1                  OK             yes
1:DIMM:12:10                 OK             yes
1:Data:12                    OK             yes
1:Data:13                    OK             yes
1:Data:8                     OK             yes
1:Data_Reduction:12          OK             yes
1:Data_Reduction:13          OK             yes
1:Data_Reduction:8           OK             yes
1:FC_Port:12:1               OK             yes
1:FC_Port:12:2               OK             yes
1:Fan:12:1                   OK             yes
1:Fan:12:2                   OK             yes
1:Flash_BBU:4:1              OK             yes
1:Flash_BBU:4:2              OK             yes
1:Flash_Canister:4:1         OK             yes
1:Flash_Canister:4:2         OK             yes
1:Flash_Card:4:1             Failed         no
1:Flash_Enclosure:4          OK             yes
1:Flash_Fan:4:1              OK             yes
1:Flash_Fan:4:2              OK             yes
1:Flash_Fan:4:3              OK             yes
1:Flash_Fan:4:4              OK             yes
1:Flash_IB_Adapter:4:1       OK             yes
1:Flash_IB_Adapter:4:2       OK             yes
1:Flash_IB_Adapter:4:3       OK             yes
1:Flash_IB_Adapter:4:4       OK             yes
1:Flash_LED_Card:4:1         OK             yes
1:Flash_PIB:4:1              OK             yes
1:Flash_PSU:4:1              OK             yes
1:Flash_PSU:4:2              OK             yes
1:IB_FlashSystem_Port:4:1    OK             yes
1:IB_FlashSystem_Port:4:3    OK             yes
1:IB_FlashSystem_Port:4:5    OK             yes
1:IB_FlashSystem_Port:4:7    OK             yes
1:IB_Module_Port:12:1        OK             yes
```

**Output:**

```
Cont:

1:IB_Module_Port:12:2       Failed         no
1:IB_Module_Port:13:1       OK             yes
1:IB_Module_Port:13:2       Failed         no
1:IB_Module_Port:8:1        OK             yes
1:IB_Module_Port:8:2        Failed         no
1:IB_Switch:1               OK             yes
1:IB_Switch:2               OK             yes
1:IB_Switch_BBU:1:1         Initializing   yes
1:IB_Switch_BBU:1:2         Initializing   yes
1:IB_Switch_BBU:2:1         Initializing   yes
1:IB_Switch_BBU:2:2         Initializing   yes
1:IB_Switch_Fan:1:1         Initializing   yes
1:IB_Switch_PSU:1:1         Initializing   yes
1:IB_Switch_PSU:1:2         Initializing   yes
1:IB_Switch_PSU:2:1         Initializing   yes
1:IB_Switch_PSU:2:2         Initializing   yes
1:IB_Switch_Port:1:12       OK             yes
1:IB_Switch_Port:1:13       OK             yes
1:IB_Switch_Port:1:20       OK             yes
1:Interface:12              OK             yes
1:Interface:13              OK             yes
1:Interface:8               OK             yes
1:Module:12                 OK             yes
1:Module:13                 OK             yes
1:Module:8                  OK             yes
1:NIC:12:1                  OK             yes
1:NIC:12:2                  OK             yes
1:PSU:12:1                  OK             yes
1:PSU:12:2                  OK             yes
1:PSU:13:1                  OK             yes
1:PSU:13:2                  OK             yes
1:PSU:8:1                   OK             yes
1:PSU:8:2                   OK             yes
1:Remote:12                 OK             yes
1:Remote:13                 OK             yes
1:Remote:8                  OK             yes
1:Vault_Device:12:1         OK             yes
1:Vault_Device:12:2         OK             yes
1:Vault_Device:13:1         OK             yes
1:Vault_Device:13:2         OK             yes
1:Vault_Device:8:1          OK             yes
1:Vault_Device:8:2          OK             yes
```

| Field ID | Field output | Default position |
|---|---|---|
| **component_id** | Component ID | 1 |
| **status** | Status | 2 |
| **currently_functioning** | Currently Functioning | 3 |
| **requires_service** | Requires Service | N/A |
| **service_reason** | Service Reason | N/A |

## Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Allowed |
| Security administrator | Disallowed |
| Read-only users | Allowed |
| Technicians | Allowed |

# Listing module configuration

Use the **module_list** command to list the configuration of all or specified modules.

```
module_list [ module=ModuleNumber ]
```

## Parameters

| Name | Description | Mandatory | Default |
|---|---|---|---|
| **module** | Lists the configuration of the specified module. | N | All modules |

The following modules are available in FlashSystem A9000 and FlashSystem A9000R:

- BBU, also referred to or known as battery module
- Boot media, also referred to or known as hard disk drive (HDD)
- Compute module, also referred to or known as grid controller or module
- IB Switch, also referred to or known as InfiniBand switch
- Module, also referred to or known as grid controller or compute module.

This command lists the following information for each module:

- Generic component status
- Module type
- Number of boot media
- Number of vault devices
- Number of FC ports
- Number of Ethernet ports for iSCSI

Additional information is available through running `module_list -t all` :

- Serial
- Original serial
- Part number
- Original part number

**Note:**

The temperature values are indicated in Celsius.

**Example:**

```
module_list
```

**Output:**

```
Component ID   Status   Currently Functioning   Target Status   Type
-------------- -------- ----------------------- --------------- ------------------------
1:Module:12    OK       yes                                     g4.0_compute_enclosure
1:Module:13    OK       yes                                     g4.0_compute_enclosure
1:Module:8     OK       yes                                     g4.0_compute_enclosure

Cont.:
Boot Media Disks   Vault Devices   FC Ports   iSCSI Ports   Temperature
------------------ --------------- ---------- ------------- -------------
2                  2               4          2             22
2                  2               4          2             22
2                  2               4          2             22
```

| Field ID | Field output | Default position |
|---|---|---|
| **component_id** | Component ID | 1 |
| **status** | Status | 2 |
| **currently_functioning** | Currently Functioning | 3 |
| **target_status** | Target Status | 4 |
| **type** | Type | 5 |
| **boot_media_disks** | Boot Media Disks | 6 |
| **vault_devices** | Vault Devices | 7 |
| **fc_port_count** | FC Ports | 8 |
| **ethernet_port_count** | iSCSI Ports | 9 |
| **temperature** | Temperature | 10 |
| **enclosure_id** | Enclosure ID | 11 |
| **avg_power** | Avg Power | N/A |
| **serial** | Serial | N/A |
| **original_serial** | Original Serial | N/A |
| **part_number** | Part Number | N/A |
| **original_part_number** | Original Part Number | N/A |
| **sas_version** | SAS | N/A |
| **infiniband_hca_version.0** | InfiniBand HCA 1 | N/A |
| **infiniband_hca_version.1** | InfiniBand HCA 2 | N/A |
| **cna_version.0** | CNA 1 | N/A |
| **cna_version.1** | CNA 2 | N/A |
| **compression_adapter_firmware.0** | Compression Adapter 1 | N/A |
| **compression_adapter_firmware.1** | Compression Adapter 2 | N/A |
| **requires_service** | Requires Service | N/A |
| **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 |

| Field ID | Field output | Default position |
|---|---|---|
| **dsa_version** | DSA Version | N/A |
| **me_version** | ME Version | N/A |
| **mcu_version** | MCU Version | N/A |
| **board_serial** | Board Serial | N/A |
| **board_part_number** | Board Part Number | N/A |
| **board_mfg_date** | Board MFG Date | N/A |
| **dasd_board_serial** | Backplane Serial | N/A |
| **dasd_board_part_number** | Backplane Part Number | N/A |
| **dasd_board_manufacturer** | Backplane Manufacturer | N/A |
| **dasd_board_mfg_date** | Backplane MFG Date | N/A |

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Allowed |
| Application administrator | Allowed |
| Security administrator | Disallowed |
| Read-only users | Allowed |
| Technicians | Allowed |

## Listing the internal temperature of modules

Use the **module_temperature_list** command to list the modules' internal temperatures in the storage system.

```
module_temperature_list [ module=ModuleNumber ]
```

### Parameters

| Name | Description | Mandatory | Default |
|---|---|---|---|
| **module** | Limits the listing to a specific module. | N | All temperatures in all modules. |

**Note:**

The temperature values are indicated in Celsius.

**Example:**

```
module_temperature_list -f all
```

**Output:**

```
Module        Ambient   PCH   RS1   RS2   MID1   MID2   RAID   DIMM AB VR   DIMM CD VR
------------- --------- ----- ----- ----- ------ ------ ------ ------------ ------------
1:Module:1    16        44    32    38    31     31     34     27           27
1:Module:11   16        46    35    40    34     33     35     26           28
1:Module:2    16        43    34    38    33     32     33     26           27
1:Module:4    16        43    34    37    33     32     34     25           26

cont:
DIMM EF VR   DIMM GH VR   CPU1   CPU2   InfiniBand HCA   PSU FR   Fibre Channel Port fc-0
------------ ------------ ------ ------ ---------------- -------- -------------------------
28           27           52     55     53               28       0
27           32           58     68     56               32       0
28           27           51     53     53               28       0
29           30           51     55     53               29       0

cont:
Fibre Channel Port fc-2    CNA=0   CNA-1   Boot Device 0   Boot Device 1   Vault Device 0
-------------------------  ------- ------- --------------- --------------- ----------------
0                          50      46      23              21              20
0                          51      52      27              25              21
0                          50      49      22              22              20
0                          48      52      24              23              20

cont:
Vault Device 1   BBU1   BBU2   CPU1 VR   CPU2 VR   Fibre Channel Port fc-1   Fibre Channel
                                                                            Port fc-3
---------------- ------ ------ --------- --------- ------------------------ -------------
21               18     18     35        36        0                        0
21               17     16     37        40        0                        0
20               18     17     35        37        0                        0
20               17     17     35        37        0                        0
```

| Field ID | Field output | Default position |
|---|---|---|
| **component_id** | Module | 1 |
| **sdr_temperatures.0** | Ambient | 2 |
| **sdr_temperatures.1** | PCH | 3 |
| **sdr_temperatures.2** | RS1 | 4 |
| **sdr_temperatures.3** | RS2 | 5 |
| **sdr_temperatures.4** | MID1 | 6 |
| **sdr_temperatures.5** | MID2 | 7 |
| **sdr_temperatures.6** | RAID | 8 |
| **sdr_temperatures.8** | CPU1 VR | N/A |
| **sdr_temperatures.9** | CPU2 VR | N/A |
| **sdr_temperatures.10** | DIMM AB VR | 9 |
| **sdr_temperatures.11** | DIMM CD VR | 10 |
| **sdr_temperatures.12** | DIMM EF VR | 11 |
| **sdr_temperatures.13** | DIMM GH VR | 12 |
| **sdr_temperatures.14** | CPU1 | 13 |
| **sdr_temperatures.15** | CPU2 | 14 |
| **ib_hca_temperature** | InfiniBand HCA | 15 |
| **sdr_temperatures.7** | PSU FR | 16 |
| **fc_adapter_temperature.0** | Fibre Channel Port fc-0 | 17 |
| **fc_adapter_temperature.1** | Fibre Channel Port fc-1 | N/A |
| **fc_adapter_temperature.2** | Fibre Channel Port fc-2 | 18 |
| **fc_adapter_temperature.3** | Fibre Channel Port fc-3 | N/A |

| Field ID | Field output | Default position |
|---|---|---|
| **cna_temperature.0** | CNA=0 | 19 |
| **cna_temperature.1** | CNA-1 | 20 |
| **boot_device_temperature.0** | Boot Device 0 | 21 |
| **boot_device_temperature.1** | Boot Device 1 | 22 |
| **vault_device_temperature.0** | Vault Device 0 | 23 |
| **vault_device_temperature.1** | Vault Device 1 | 24 |
| **bbu_temperature.0** | BBU1 | 25 |
| **bbu_temperature.1** | BBU2 | 26 |

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Allowed |
| Security administrator | Disallowed |
| Read-only users | Allowed |
| Technicians | Allowed |

## Listing boot media devices in the system

Use the **boot_media_list** command to lists boot media devices in the storage system.

```
boot_media_list [ module=ModuleNumber | boot_media=BootMediaDevice ]
```

### Parameters

| Name | Description | Mandatory | Default |
|---|---|---|---|
| **module** | Limits the listing to the specific module. | N | Boot media devices in a specific module. |
| **boot_media** | Limits the listing to a specific boot media. | N | A specific boot media device. |

**Note:**

The temperature values are indicated in Celsius.

**Example:**

```
boot_media_list -f all
```

**Output:**

```
Component ID        Status   Currently Functioning   Hardware Status   Vendor
------------------- -------- ---------------------- ----------------- --------
1:Boot_Media:11:1   OK       yes                    OK                IBM-ESXS
1:Boot_Media:11:2   OK       yes                    OK                IBM-ESXS
1:Boot_Media:1:1    OK       yes                    OK                IBM-ESXS
1:Boot_Media:1:2    OK       yes                    OK                IBM-ESXS
1:Boot_Media:2:1    OK       yes                    OK                IBM-ESXS
1:Boot_Media:2:2    OK       yes                    OK                IBM-ESXS
1:Boot_Media:4:1    OK       yes                    OK                IBM-ESXS
1:Boot_Media:4:2    OK       yes                    OK                IBM-ESXS

Model              Serial   FW     Temperature   Original Serial   Part #
------------------ -------- ------ ------------- ----------------- --------
HUC101860CS420     03V0E75K J5H2   26            03V0E75K
HUC101860CS420     03V0LPEK J5H2   25            03V0LPEK
HUC101860CS420     03V0X90H J5H2   21            03V0X90H
HUC101860CS420     03V0AHME J5H2   20            03V0AHME
HUC101860CS420     03V117XE J5H2   22            03V117XE
HUC101860CS420     03V0DAGK J5H2   21            03V0DAGK
HUC101860CS420     03V0DJ4K J5H2   23            03V0DJ4K
HUC101860CS420     03V0H38K J5H2   22            03V0H38K

Original Part #  Size (GB)   Requires Service   Service Reason
---------------- ---------- ------------------ ----------------
                 600 GB
                 600 GB
                 600 GB
                 600 GB
                 600 GB
                 600 GB
                 600 GB
                 600 GB

Rebuild Progress   Rebuild Time (sec.)
------------------ ---------------------
N/A                N/A
N/A                N/A
N/A                N/A
N/A                N/A
N/A                N/A
N/A                N/A
N/A                N/A
N/A                N/A
```

| Field ID | Field output | Default position |
|---|---|---|
| **component_id** | Component ID | 1 |
| **status** | Status | 2 |
| **currently_functioning** | Currently Functioning | 3 |
| **hardware_status** | Hardware Status | 4 |
| **vendor** | Vendor | 5 |
| **model** | Model | 6 |
| **serial** | Serial | 7 |
| **fw_revision** | FW | 8 |
| **temperature** | Temperature | N/A |
| **original_serial** | Original Serial | N/A |
| **part_number** | Part # | N/A |
| **original_part_number** | Original Part # | N/A |
| **fru_pn** | FRU PN | N/A |
| **original_fru_pn** | Original FRU PN | N/A |
| **size** | Size | N/A |
| **requires_service** | Requires Service | N/A |

| Field ID | Field output | Default position |
|---|---|---|
| **service_reason** | Service Reason | N/A |
| **rebuild_progress** | Rebuild Progress | N/A |
| **rebuild_time** | Rebuild Time | N/A |
| **media_error_count** | Media Errors | N/A |
| **other_error_count** | Other Errors | N/A |
| **predictive_failure_count** | Predictive Failures | N/A |

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Allowed |
| Security administrator | Disallowed |
| Read-only users | Allowed |
| Technicians | Allowed |

## Listing vault devices in the system

Use the **vault_device_list** command to lists special vault device statuses.

```
vault_device_list [ module=ModuleNumber  | vault_device=ComponentId ]
```

### Parameters

| Name | Description | Mandatory | Default |
|---|---|---|---|
| **module** | Limits the listing to a specific module. | N | All vault devices in all modules. |
| **vault_device** | Vault devices for which special statuses are to be listed. | N | All vault devices. |

**Note:**

The temperature values are indicated in Celsius.

**Example:**

```
vault_device_list
```

**Output:**

```
Component ID          Status   Currently Functioning   Capacity   Target Status   Vendor
--------------------- -------- ---------------------- ---------- --------------- --------
1:Vault_Device:14:1   OK       yes                    250GB                      LENOVO-X
1:Vault_Device:14:2   OK       yes                    250GB                      LENOVO-X
1:Vault_Device:3:1    OK       yes                    250GB                      LENOVO-X
1:Vault_Device:3:2    OK       yes                    250GB                      LENOVO-X
1:Vault_Device:5:1    OK       yes                    250GB                      LENOVO-X
1:Vault_Device:5:2    OK       yes                    250GB                      LENOVO-X
1:Vault_Device:6:1    OK       yes                    250GB                      LENOVO-X
1:Vault_Device:6:2    OK       yes                    250GB                      LENOVO-X

Cont.:
Model             Serial     Firmware   FRU       Temperature   Encryption State
----------------- ---------- ---------- --------- ------------- ------------------
HUSMR1625ASS20E   0PVGJTPA   P4C9       00NA685   22            Ready
HUSMR1625ASS20E   0PVGHN6A   P4C9       00NA685   22            Ready
HUSMR1625ASS20E   0PVJ2PEA   P4C9       00NA685   21            Ready
HUSMR1625ASS20E   0PVJ251A   P4C9       00NA685   21            Ready
HUSMR1625ASS20E   0PVJ99UA   P4C9       00NA685   20            Ready
HUSMR1625ASS20E   0PVJRB5A   P4C9       00NA685   21            Ready
HUSMR1625ASS20E   0PVJS39A   P4C9       00NA685   20            Ready
HUSMR1625ASS20E   0PVJ9RAA   P4C9       00NA685   21            Ready
```

| Field ID | Field output | Default position |
|---|---|---|
| **component_id** | Component ID | 1 |
| **status** | Status | 2 |
| **currently_functioning** | Currently Functioning | 3 |
| **capacity_in_bytes** | Capacity | N/A |
| **capacity** | Capacity | 4 |
| **target_status** | Target Status | 5 |
| **vendor** | Vendor | 6 |
| **original_vendor** | Original Vendor | N/A |
| **model** | Model | 7 |
| **original_model** | Original Model | N/A |
| **serial** | Serial | 8 |
| **original_serial** | Original Serial | N/A |
| **firmware** | Firmware | 9 |
| **original_firmware** | Original Firmware | N/A |
| **part_number** | FRU | 10 |
| **original_part_number** | Original FRU | N/A |
| **temperature** | Temperature | 11 |
| **encryption_state** | Encryption State | 12 |
| **hw_mon_node_id** | Hw Node Owner | N/A |
| **requires_service** | Requires Service | N/A |
| **service_reason** | Service Reason | N/A |
| **revision** | Revision | N/A |
| **drive_pn** | Drive P/N | N/A |
| **original_drive_pn** | Original Drive P/N | N/A |
| **fru_pn** | FRU P/N | N/A |
| **original_fru_pn** | Original FRU P/N | N/A |
| **desc.bgd_scan** | Background Scan | N/A |

| Field ID | Field output | Default position |
|---|---|---|
| **desc.disk_id** | Device ID | N/A |
| **desc.last_sample_serial** | Last Sample Serial | N/A |
| **desc.last_sample_time** | Last Sample Time | N/A |
| **desc.power_is_on** | Power On | N/A |
| **desc.power_on_hours** | Power On Hours | N/A |
| **desc.power_on_minutes** | Power On Minutes | N/A |
| **desc.last_time_pom_was_mod** | Last Time Power On Minutes Was Modified | N/A |
| **desc.read_fail** | Read Fail | N/A |
| **desc.smart_code** | SMART Code | N/A |
| **desc.smart_fail** | SMART Fail | N/A |
| **desc.temperature_status.reported_severity** | Reported Temperature Severity | N/A |
| **desc.temperature_status.reported_temperature** | Reported Temperature | N/A |
| **desc.temperature_status.temperature** | Device Temperature | N/A |
| **desc.sw_encryption_active** | Software-Based Encryption Active | N/A |

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Allowed |
| Security administrator | Allowed |
| Read-only users | Allowed |
| Technicians | Allowed |

## Listing BBUs in the system

Use the **module_bbu_list** command to list the BBUs in the storage system modules.

```
module_bbu_list [ module=ModuleNumber | module_bbu=BbuNumber ]
```

### Parameters

| Name | Description | Mandatory | Default |
|---|---|---|---|
| **module** | Limits the listing to a specific module. | N | All BBUs in all modules. |
| **module_bbu** | Limits the listing to a specific bbu. | N | ALL BBUs. |

**Example:**

```
module_bbu_list -f all
```

**Output:**

```
Component ID   Status   Curr Functioning   State   Hardware Status   Remaining Capacity
-------------- -------- ------------------ ------- ----------------- --------------------
1:BBU:14:1     OK       yes                Full    OK                891
1:BBU:14:2     OK       yes                Full    OK                877
1:BBU:3:1      OK       yes                Full    OK                787
1:BBU:3:2      OK       yes                Full    OK                860
1:BBU:5:1      OK       yes                Full    OK                792
1:BBU:5:2      OK       yes                Full    OK                898
1:BBU:6:1      OK       yes                Full    OK                817
1:BBU:6:2      OK       yes                Full    OK                814

cont:
Full Charge Capacity   Charged %   Time to Empty   Time to Full   Charger State
---------------------- ----------- --------------- -------------- ---------------
891                    100         1600200         0              in progress
877                    100         0               0              in progress
787                    100         0               0              in progress
860                    100         0               0              in progress
792                    100         0               0              in progress
898                    100         0               0              in progress
817                    100         2944800         0              in progress
814                    100         0               0              in progress

cont:
Calibration State   Calibration Time
------------------- ------------------
Idle                0
Idle                0
Idle                0
Idle                0
Idle                0
Idle                0
Idle                0
Idle                0
```

| Field ID | Field output | Default position |
|---|---|---|
| **component_id** | Component ID | 1 |
| **status** | Status | 2 |
| **currently_functioning** | Curr Functioning | 3 |
| **required_service** | Requires Service | N/A |
| **service_reason** | Service Reason | N/A |
| **bbu_state** | State | 4 |
| **bbu_status** | Hardware Status | 5 |
| **remaining_capacity** | Remaining Capacity | 6 |
| **full_charge_capacity** | Full Charge Capacity | 7 |
| **percent_charged** | Charged % | 8 |
| **time_to_empty** | Time to Empty | 9 |
| **time_to_full** | Time to Full | 10 |
| **charger_state** | Charger State | 11 |
| **calib_state** | Calibration State | 12 |
| **calib_requested** | Calibration Needed | N/A |
| **last_succ_calib_date** | Successful Calibration Time | N/A |
| **last_calib_date** | Calibration Time | 13 |
| **last_calib_result** | Calibration Result | N/A |
| **insertion_date** | Inserted | N/A |
| **manuf_date** | Manufactured | N/A |

| Field ID | Field output | Default position |
|---|---|---|
| **fw** | FW Version | N/A |
| **epow_cable_present** | EPOW Cable Present | N/A |
| **power_sense_cable_present** | Power Sense Cable Present | N/A |
| **epow_simulate** | EPOW Simulate | N/A |
| **epow_asserted** | EPOW Asserted | N/A |
| **cycle_count** | Cycles | N/A |
| **temperature_tenths_celsius** | Temp /10C | N/A |
| **charger_enabled** | Charger Enabled | N/A |
| **slow_charge_enabled** | Slow Charge Enabled | N/A |
| **discharge_enabled** | Discharge Enabled | N/A |
| **ps2_present** | PS2 Present | N/A |
| **charge_now** | Nominal Available Capacity mAh | N/A |
| **voltage_now** | Voltage Now mV | N/A |
| **current_now** | Current Now mA | N/A |
| **power_avg** | Power Average mW | N/A |
| **charge_full** | Full Available Capacity mAh | N/A |
| **charge_full_design** | Design Charge | N/A |
| **energy_now** | Energy now mWh | N/A |
| **at_rate** | At Rate | N/A |
| **at_rate_tte** | At Rate Time to Empty | N/A |
| **charge_now_sufficient** | Charge Now Sufficient | N/A |
| **endurance_start_monotonic_time** | Endurance Start Monotonic Time | N/A |
| **serial** | Serial | N/A |
| **original_serial** | Original Serial | N/A |
| **part_number** | Part # | N/A |
| **original_part_number** | Original Part # | N/A |
| **fru** | FRU | N/A |
| **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 |

| User Category | Permission |
|---|---|
| Storage integration administrator | Allowed |
| Application administrator | Allowed |
| Security administrator | Disallowed |
| Read-only users | Allowed |
| Technicians | Allowed |

# Listing PSUs in the system

Use the **module_psu_list** command to list PSUs in the system.

```
module_psu_list [ module=ModuleNumber  | psu=ComponentId ]
```

## Parameters

| Name | Description | Mandatory | Default |
|---|---|---|---|
| **module** | Limits the listing to a specific module. | N | All PSUs in all modules. |
| **psu** | Lists only a specific PSU. | N | A specific PSU. |

**Example:**

```
module_psu_list -f all
```

**Output:**

```
Component ID   Status   Currently Functioning   Location
-------------- -------- ---------------------- ---------------
1:PSU:1:1      Failed   no                     Power Supply 1
1:PSU:1:2      OK       yes                    Power Supply 2
1:PSU:2:1      OK       yes                    Power Supply 1
1:PSU:2:2      Failed   no                     Power Supply 2
1:PSU:4:1      OK       yes                    Power Supply 1
1:PSU:4:2      Failed   no                     Power Supply 2

cont:
Sensor status                             Serial number   Part number
----------------------------------------- --------------- -------------
Presence NOT detected                     N/A             N/A
Presence detected                         K115148J01L     94Y8143
Presence detected                         K115148B06E     94Y8143
Presence detected, Power Supply AC lost   K115148B01P     94Y8143
Presence detected                         K115148B0AP     94Y8143
Presence detected, Power Supply AC lost   K115148J01B     94Y8143

cont:
Requires Service   Service Reason
------------------ -------------------------
REPLACE            MODULE_PSU__NOT_DETECTED


COMPONENT_TEST     MODULE_PSU__BAD_POWER_INPUT

COMPONENT_TEST     MODULE_PSU__BAD_POWER_INPUT
```

| Field ID | Field output | Default position |
|---|---|---|
| **component_id** | Component ID | 1 |
| **status** | Status | 2 |
| **currently_functioning** | Currently Functioning | 3 |

| Field ID | Field output | Default position |
|---|---|---|
| **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 |
| **last_valid_serial** | Last Valid Serial number | N/A |
| **part_number** | Part number | N/A |
| **manufacturer** | Manufacturer | N/A |
| **mfg_date** | Manufacturing Date | N/A |

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Allowed |
| Security administrator | Disallowed |
| Read-only users | Allowed |
| Technicians | Allowed |

## Listing compression adapters in the system

Use the **compression_adapter_list** command to list compression adapters in the system.

```
compression_adapter_list [ module=ModuleNumber  | compression_adapter=ComponentId ]
```

### Parameters

| Name | Description | Mandatory | Default |
|---|---|---|---|
| **module** | Limits the listing to a specific module. | N | All compression adapters in all modules. |
| **compression_adapter** | Lists only a specific compression adapter. | N | A specific compression adapter. |

**Note:** The Technician Assistant tool and XCLI Utility refer to the Ethernet adapter as compression adapter.

Compression adapters are used to increase the speed of I/O transfers to and from compressed volumes.

**Example:**

```
compression_adapter_list
```

**Output:**

```
Component ID                Status   Currently Functioning   Firmware   Hardware
--------------------------- -------- ---------------------- ---------- ----------
1:Compression_Adapter:6:1   OK       yes                    1.0.12     A0 SKU3
1:Compression_Adapter:6:2   OK       yes                    1.0.12     A0 SKU3

Driver
---------------
1.0.12 3d60d12
1.0.12 3d60d12
```

| Field ID | Field output | Default position |
|---|---|---|
| **component_id** | Component ID | 1 |
| **status** | Status | 2 |
| **currently_functioning** | Currently Functioning | 3 |
| **serial** | Serial | 4 |
| **firmware_version** | Firmware | 5 |
| **hardware_version** | Hardware | 6 |
| **driver_version** | Driver | 7 |
| **type** | Type | N/A |
| **mmp_version** | MMP | N/A |
| **pci_address** | PCI Address | N/A |
| **quick_assist_api_cy_version** | QuickAssist API CY | N/A |
| **quick_assist_api_dc_version** | QuickAssist API DC | N/A |
| **threading_mode** | Threading Mode | N/A |
| **requires_service** | Requires Service | N/A |
| **service_reason** | Service Reason | N/A |

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Allowed |
| Security administrator | Disallowed |
| Read-only users | Allowed |
| Technicians | Allowed |

## Listing fans in a module

Use the **fan_list** command to list fans in the specified module.

```
fan_list [ module=ModuleNumber  | fan=ComponentId ]
```

### Parameters

| Name | Description | Mandatory | Default |
|---|---|---|---|
| **module** | Limits the listing to a specific module. | N | All fans in all modules. |

| Name | Description | Mandatory | Default |
|---|---|---|---|
| **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 |

| Field ID | Field output | Default position |
|---|---|---|
| **max_rpm_a** | Max Speed | N/A |
| **location_b** | Peer Location | N/A |
| **rpm_b** | Peer Speed | 5 |
| **rpm_b** | Peer Speed | N/A |
| **min_rpm_b** | Peer Min Speed | N/A |
| **max_rpm_b** | Peer Max Speed | N/A |
| **requires_service** | Requires Service | N/A |
| **service_reason** | Service Reason | N/A |

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Allowed |
| Security administrator | Disallowed |
| Read-only users | Allowed |
| Technicians | Allowed |

## Listing NICs in the system

Use the **nic_list** command to list the NICs in the storage system.

```
nic_list [ module=ModuleNumber  | nic=ComponentId ]
```

### Parameters

| Name | Description | Mandatory | Default |
|---|---|---|---|
| **module** | Limits the listing to a specific module. | N | All NICs in all modules. |
| **nic** | Lists only a specific NIC. | N | A specific NIC. |

**Example:**

```
nic_list -f all
```

**Output:**

```
Component ID   Status   Currently Functioning   Hardware Status   Device Name
-------------- -------- ----------------------- ----------------- -------------
1:NIC:10:1     OK       yes                     OK                eth0
1:NIC:10:2     OK       yes                     OK                eth1
1:NIC:10:3     OK       yes                     OK                eth2
1:NIC:10:4     OK       yes                     OK                eth3
1:NIC:10:5     OK       yes                     OK                eth4
1:NIC:10:6     OK       yes                     OK                eth5
1:NIC:11:1     OK       yes                     OK                eth0
1:NIC:11:2     OK       yes                     OK                eth1
1:NIC:11:3     OK       yes                     OK                eth2
1:NIC:11:4     OK       yes                     OK                eth3
1:NIC:11:5     OK       yes                     OK                eth4
1:NIC:11:6     OK       yes                     OK                eth5
1:NIC:7:1      OK       yes                     OK                eth0
1:NIC:7:2      OK       yes                     OK                eth1
1:NIC:7:3      OK       yes                     OK                eth2
1:NIC:7:4      OK       yes                     OK                eth3
1:NIC:9:1      OK       yes                     OK                eth0
1:NIC:9:2      OK       yes                     OK                eth1
1:NIC:9:3      OK       yes                     OK                eth2
1:NIC:9:4      OK       yes                     OK                eth3
1:NIC:9:5      OK       yes                     OK                eth4
1:NIC:9:6      OK       yes                     OK                eth5

Cont.:

Serial              Original Serial     Part #
------------------- ------------------- ---------------------------------------
40:f2:e9:af:26:b0   40:f2:e9:af:26:b0   14e4_1657_40f2e9af26b0_5719-v1.38
40:f2:e9:af:26:b1   40:f2:e9:af:26:b1   14e4_1657_40f2e9af26b1_5719-v1.38
40:f2:e9:af:26:b2   40:f2:e9:af:26:b2   14e4_1657_40f2e9af26b2_5719-v1.38
40:f2:e9:af:26:b3   40:f2:e9:af:26:b3   14e4_1657_40f2e9af26b3_5719-v1.38
f4:52:14:6e:8f:70   f4:52:14:6e:8f:70   15b3_1007_708f6e00031452f4_2.35.5100
f4:52:14:6e:8f:71   f4:52:14:6e:8f:71   15b3_1007_708f6e00031452f4_2.35.5100
40:f2:e9:af:24:48   40:f2:e9:af:24:48   14e4_1657_40f2e9af2448_5719-v1.38
40:f2:e9:af:24:49   40:f2:e9:af:24:49   14e4_1657_40f2e9af2449_5719-v1.38
40:f2:e9:af:24:4a   40:f2:e9:af:24:4a   14e4_1657_40f2e9af244a_5719-v1.38
40:f2:e9:af:24:4b   40:f2:e9:af:24:4b   14e4_1657_40f2e9af244b_5719-v1.38
f4:52:14:6e:8d:30   f4:52:14:6e:8d:30   15b3_1007_308d6e00031452f4_2.35.5100
f4:52:14:6e:8d:31   f4:52:14:6e:8d:31   15b3_1007_308d6e00031452f4_2.35.5100
40:f2:e9:af:2a:90   40:f2:e9:af:2a:90   14e4_1657_40f2e9af2a90_5719-v1.38
40:f2:e9:af:2a:91   40:f2:e9:af:2a:91   14e4_1657_40f2e9af2a91_5719-v1.38
40:f2:e9:af:2a:92   40:f2:e9:af:2a:92   14e4_1657_40f2e9af2a92_5719-v1.38
40:f2:e9:af:2a:93   40:f2:e9:af:2a:93   14e4_1657_40f2e9af2a93_5719-v1.38
40:f2:e9:af:23:b8   40:f2:e9:af:23:b8   14e4_1657_40f2e9af23b8_5719-v1.38
40:f2:e9:af:23:b9   40:f2:e9:af:23:b9   14e4_1657_40f2e9af23b9_5719-v1.38
40:f2:e9:af:23:ba   40:f2:e9:af:23:ba   14e4_1657_40f2e9af23ba_5719-v1.38
40:f2:e9:af:23:bb   40:f2:e9:af:23:bb   14e4_1657_40f2e9af23bb_5719-v1.38
f4:52:14:6e:8f:20   f4:52:14:6e:8f:20   15b3_1007_208f6e00031452f4_2.35.5100
f4:52:14:6e:8f:21   f4:52:14:6e:8f:21   15b3_1007_208f6e00031452f4_2.35.5100

Cont.:

Original Part Number                   Requires Service   Service Reason
-------------------------------------- ------------------ ----------------
14e4_1657_40f2e9af26b0_5719-v1.38
14e4_1657_40f2e9af26b1_5719-v1.38
14e4_1657_40f2e9af26b2_5719-v1.38
14e4_1657_40f2e9af26b3_5719-v1.38
15b3_1007_708f6e00031452f4_2.35.5100
15b3_1007_708f6e00031452f4_2.35.5100
14e4_1657_40f2e9af2448_5719-v1.38
14e4_1657_40f2e9af2449_5719-v1.38
14e4_1657_40f2e9af244a_5719-v1.38
14e4_1657_40f2e9af244b_5719-v1.38
15b3_1007_308d6e00031452f4_2.35.5100
15b3_1007_308d6e00031452f4_2.35.5100
14e4_1657_40f2e9af2a90_5719-v1.38
14e4_1657_40f2e9af2a91_5719-v1.38
14e4_1657_40f2e9af2a92_5719-v1.38
14e4_1657_40f2e9af2a93_5719-v1.38
14e4_1657_40f2e9af23b8_5719-v1.38
14e4_1657_40f2e9af23b9_5719-v1.38
14e4_1657_40f2e9af23ba_5719-v1.38
14e4_1657_40f2e9af23bb_5719-v1.38
15b3_1007_208f6e00031452f4_2.35.5100
15b3_1007_208f6e00031452f4_2.35.5100
```

| Field ID | Field output | Default position |
|---|---|---|
| **component_id** | Component ID | 1 |
| **status** | Status | 2 |
| **currently_functioning** | Currently Functioning | 3 |
| **hardware_status** | Hardware Status | 4 |
| **device_name** | Device Name | 5 |
| **serial** | Serial | N/A |
| **original_serial** | Original Serial | N/A |
| **part_number** | Part # | N/A |
| **original_part_number** | Original Part Number | N/A |
| **requires_service** | Requires Service | N/A |
| **service_reason** | Service Reason | N/A |

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Allowed |
| Security administrator | Disallowed |
| Read-only users | Allowed |
| Technicians | Allowed |

## Listing DIMMs in the modules

Use the **dimm_list** command to list the DIMMs in the modules.

```
dimm_list [ module=ModuleNumber  | dimm=ComponentId ]
```

### Parameters

| Name | Description | Mandatory | Default |
|---|---|---|---|
| **module** | Limits the listing to a specific module. | N | All DIMMs in all modules. |
| **dimm** | Lists only a specific DIMM. | N | A specific DIMM. |

The memory modules (DIMMs) run the microcode and the data cache in the grid controller.

**Example:**

```
dimm_list -f all
```

**Output:**

```
Component ID   Status   Currently Functioning   Hardware Status   DIMM Id   CPU
-------------- -------- ----------------------- ----------------- --------- -----
1:DIMM:7:1     OK       yes                     OK                1         1
1:DIMM:7:10    OK       yes                     OK                10        1
1:DIMM:7:11    OK       yes                     OK                11        1
1:DIMM:7:12    OK       yes                     OK                12        1
1:DIMM:7:13    OK       yes                     OK                13        2
1:DIMM:7:14    OK       yes                     OK                14        2
1:DIMM:7:15    OK       yes                     OK                15        2
1:DIMM:7:16    OK       yes                     OK                16        2
1:DIMM:7:17    OK       yes                     OK                17        2
1:DIMM:7:18    OK       yes                     OK                18        2
1:DIMM:7:19    OK       yes                     OK                19        2
1:DIMM:7:2     OK       yes                     OK                2         1
1:DIMM:7:20    OK       yes                     OK                20        2
1:DIMM:7:21    OK       yes                     OK                21        2
1:DIMM:7:22    OK       yes                     OK                22        2
1:DIMM:7:23    OK       yes                     OK                23        2
1:DIMM:7:24    OK       yes                     OK                24        2
1:DIMM:7:3     OK       yes                     OK                3         1
1:DIMM:7:4     OK       yes                     OK                4         1
1:DIMM:7:5     OK       yes                     OK                5         1
1:DIMM:7:6     OK       yes                     OK                6         1
1:DIMM:7:7     OK       yes                     OK                7         1
1:DIMM:7:8     OK       yes                     OK                8         1
1:DIMM:7:9     OK       yes                     OK                9         1
1:DIMM:9:1     OK       yes                     OK                1         1
1:DIMM:9:10    OK       yes                     OK                10        1
1:DIMM:9:11    OK       yes                     OK                11        1
1:DIMM:9:12    OK       yes                     OK                12        1
1:DIMM:9:13    OK       yes                     OK                13        2
1:DIMM:9:14    OK       yes                     OK                14        2
1:DIMM:9:15    OK       yes                     OK                15        2
1:DIMM:9:16    OK       yes                     OK                16        2
1:DIMM:9:17    OK       yes                     OK                17        2
1:DIMM:9:18    OK       yes                     OK                18        2
1:DIMM:9:19    OK       yes                     OK                19        2
1:DIMM:9:2     OK       yes                     OK                2         1
1:DIMM:9:20    OK       yes                     OK                20        2
1:DIMM:9:21    OK       yes                     OK                21        2
1:DIMM:9:22    OK       yes                     OK                22        2
1:DIMM:9:23    OK       yes                     OK                23        2
1:DIMM:9:24    OK       yes                     OK                24        2
1:DIMM:9:3     OK       yes                     OK                3         1
1:DIMM:9:4     OK       yes                     OK                4         1
1:DIMM:9:5     OK       yes                     OK                5         1
1:DIMM:9:6     OK       yes                     OK                6         1
1:DIMM:9:7     OK       yes                     OK                7         1
1:DIMM:9:8     OK       yes                     OK                8         1
1:DIMM:9:9     OK       yes                     OK                9         1
```

Cont.:

**Output:**

```
Size(Mb)   Speed(MHz)   Configured Clock Speed(MHz)   Manufacturer   Serial
---------- ------------ ----------------------------- -------------- ----------
16384      2133         2133                          Hynix          505F63F9
16384      2133         2133                          Samsung        39542977
16384      2133         2133                          Samsung        3953EA6C
16384      2133         2133                          Samsung        3953F39A
16384      2133         2133                          Samsung        3953F240
16384      2133         2133                          Samsung        39542562
16384      2133         2133                          Samsung        3954300B
16384      2133         2133                          Samsung        39546472
16384      2133         2133                          Samsung        39540BB2
16384      2133         2133                          Samsung        3953FB59
16384      2133         2133                          Samsung        3954074A
16384      2133         2133                          Samsung        3953F241
16384      2133         2133                          Samsung        395404E0
16384      2133         2133                          Samsung        395425D8
16384      2133         2133                          Samsung        39542BF2
16384      2133         2133                          Samsung        395426EF
16384      2133         2133                          Samsung        3953EB61
16384      2133         2133                          Samsung        39542AD0
16384      2133         2133                          Samsung        39542973
16384      2133         2133                          Samsung        39542ACF
16384      2133         2133                          Samsung        3953E982
16384      2133         2133                          Samsung        39542568
16384      2133         2133                          Samsung        3953EA4A
16384      2133         2133                          Samsung        3953E993
16384      2133         2133                          Hynix          804AC8C2
16384      2133         2133                          Hynix          707387FA
16384      2133         2133                          Hynix          7073895C
16384      2133         2133                          Hynix          3077315A
16384      2133         2133                          Hynix          90655EDF
16384      2133         2133                          Hynix          7073885E
16384      2133         2133                          Hynix          30772789
16384      2133         2133                          Hynix          707388BB
16384      2133         2133                          Hynix          90655FC6
16384      2133         2133                          Hynix          70738960
16384      2133         2133                          Hynix          70738871
16384      2133         2133                          Hynix          7073881A
16384      2133         2133                          Hynix          307C97D6
16384      2133         2133                          Hynix          7073880D
16384      2133         2133                          Hynix          70738819
16384      2133         2133                          Hynix          30773136
16384      2133         2133                          Hynix          30772FF8
16384      2133         2133                          Hynix          70738850
16384      2133         2133                          Hynix          3079B2CD
16384      2133         2133                          Hynix          307A04D2
16384      2133         2133                          Hynix          70738951
16384      2133         2133                          Hynix          70738939
16384      2133         2133                          Hynix          70738A0B
16384      2133         2133                          Hynix          70738823
```

Cont.:

**Output:**

```
Original          Serial Part #       Original Part Number
----------------- ------------------  ----------------------
505F63F9          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
39542977          M393A2G40DB0-CPB    M393A2G40DB0-CPB
3953EA6C          M393A2G40DB0-CPB    M393A2G40DB0-CPB
3953F39A          M393A2G40DB0-CPB    M393A2G40DB0-CPB
3953F240          M393A2G40DB0-CPB    M393A2G40DB0-CPB
39542562          M393A2G40DB0-CPB    M393A2G40DB0-CPB
3954300B          M393A2G40DB0-CPB    M393A2G40DB0-CPB
39546472          M393A2G40DB0-CPB    M393A2G40DB0-CPB
39540BB2          M393A2G40DB0-CPB    M393A2G40DB0-CPB
3953FB59          M393A2G40DB0-CPB    M393A2G40DB0-CPB
3954074A          M393A2G40DB0-CPB    M393A2G40DB0-CPB
3953F241          M393A2G40DB0-CPB    M393A2G40DB0-CPB
395404E0          M393A2G40DB0-CPB    M393A2G40DB0-CPB
395425D8          M393A2G40DB0-CPB    M393A2G40DB0-CPB
39542BF2          M393A2G40DB0-CPB    M393A2G40DB0-CPB
395426EF          M393A2G40DB0-CPB    M393A2G40DB0-CPB
3953EB61          M393A2G40DB0-CPB    M393A2G40DB0-CPB
39542AD0          M393A2G40DB0-CPB    M393A2G40DB0-CPB
39542973          M393A2G40DB0-CPB    M393A2G40DB0-CPB
39542ACF          M393A2G40DB0-CPB    M393A2G40DB0-CPB
3953E982          M393A2G40DB0-CPB    M393A2G40DB0-CPB
39542568          M393A2G40DB0-CPB    M393A2G40DB0-CPB
3953EA4A          M393A2G40DB0-CPB    M393A2G40DB0-CPB
3953E993          M393A2G40DB0-CPB    M393A2G40DB0-CPB
804AC8C2          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
707387FA          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
7073895C          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
3077315A          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
90655EDF          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
7073885E          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
30772789          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
707388BB          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
90655FC6          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
70738960          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
70738871          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
7073881A          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
307C97D6          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
7073880D          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
70738819          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
30773136          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
30772FF8          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
70738850          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
3079B2CD          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
307A04D2          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
70738951          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
70738939          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
70738A0B          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
70738823          HMA42GR7MFR4N-TF    HMA42GR7MFR4N-TF
```

Cont:

**Output:**

```
Requires Service   Service Reason
-----------------  ------------------
```

| Field ID | Field output | Default position |
|---|---|---|
| **component_id** | Component ID | 1 |
| **status** | Status | 2 |
| **currently_functioning** | Currently Functioning | 3 |
| **hardware_status** | Hardware Status | 4 |
| **dimm_id** | DIMM Id | N/A |
| **cpu** | CPU | N/A |
| **size** | Size(Mb) | N/A |

| 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 |
| Application administrator | Allowed |
| Security administrator | Disallowed |
| Read-only users | Allowed |
| Technicians | Allowed |

## Listing InfiniBand HCA Adapters in the storage system

Use the **hca_list** command to list the InfiniBand host card adapters (HCAs) in the storage system.

```
hca_list [ module=ModuleNumber | hca=ComponentId ]
```

### Parameters

| Name | Description | Mandatory | Default |
|---|---|---|---|
| **module** | Limits the listing to a specific module. | N | All InfiniBand HCA adapters in all modules. |
| **hca** | Lists only a specific HCA. | N | A specific InfiniBand HCA. |

**Example:**

```
hca_list
```

**Output:**

```
Component ID   Status   Currently Functioning   Board Description
-------------- -------- ----------------------- --------------------------
1:HCA:10:1     OK       yes                     CB194A - Connect-IB QSFP
1:HCA:7:1      OK       yes                     CB194A - Connect-IB QSFP
1:HCA:9:1      OK       yes                     CB194A - Connect-IB QSFP

cont:
Board ID        Part Number
--------------- -------------
MT_1210110019   46W0572
MT_1210110019   46W0572
MT_1210110019   46W0572
```

| Field ID | Field output | Default position |
|---|---|---|
| **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 |

| Field ID | Field output | Default position |
|---|---|---|
| **original_board_type** | Original Board Type | N/A |
| **serial** | Serial | N/A |
| **original_serial** | Original Serial | N/A |
| **part_number** | Part Number | 6 |
| **original_part_number** | Original Part Number | N/A |
| **hardware_revision** | Hardware Revision | N/A |
| **requires_service** | Requires Service | N/A |
| **service_reason** | Service Reason | N/A |
| **adapter_id** | HCA Id | N/A |
| **guid** | GUID | N/A |
| **vendor_part_id** | Vendor Part ID | N/A |

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Allowed |
| Security administrator | Disallowed |
| Read-only users | Allowed |
| Technicians | Allowed |

## Listing CNA adapters in the system

Use the **cna_list** command to list CNA adapters in the storage system.

```
cna_list [ module=ModuleNumber | cna=ComponentId ]
```

### Parameters

| Name | Description | Mandatory | Default |
|---|---|---|---|
| **module** | Limits the listing to a specific module. | N | All CNA adapters in all modules. |
| **cna** | Lists only a specific CNA. | N | A specific CNA. |

A converged network adapter (CNA) is a single network interface card that contains both a Fibre Channel host bus adapter and a TCP/IP Ethernet NIC. It connects servers to FC-based storage area networks (SANs) and Ethernet-based local area networks (LANs).

**Example:**

```
cna_list
```

**Output:**

```
Component ID   Status   Currently Functioning   Board Description
-------------- -------- ----------------------- ----------------------------
1:CNA:10:1     OK       yes                     CX312B - ConnectX-3 Pro SFP+
1:CNA:9:1      OK       yes                     CX312B - ConnectX-3 Pro SFP+

Board ID        Part Number
--------------- -------------
MT_1200111023   MCX312B-XCCT
MT_1200111023   MCX312B-XCCT
```

| Field ID | Field output | Default position |
|---|---|---|
| **component_id** | Component ID | 1 |
| **status** | Status | 2 |
| **currently_functioning** | Currently Functioning | 3 |
| **version** | Version | N/A |
| **board_description** | Board Description | 4 |
| **original_board_description** | Original Board Description | N/A |
| **board_id** | Board ID | 5 |
| **original_board_id** | Original Board ID | N/A |
| **board_type** | Board Type | N/A |
| **original_board_type** | Original Board Type | N/A |
| **serial** | Serial | N/A |
| **original_serial** | Original Serial | N/A |
| **part_number** | Part Number | 6 |
| **original_part_number** | Original Part Number | N/A |
| **hardware_revision** | Hardware Revision | N/A |
| **requires_service** | Requires Service | N/A |
| **service_reason** | Service Reason | N/A |
| **adapter_id** | HCA Id | N/A |
| **guid** | GUID | N/A |
| **vendor_part_id** | Vendor Part ID | N/A |

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Allowed |
| Security administrator | Disallowed |
| Read-only users | Allowed |
| Technicians | Allowed |

## Listing module LEDs in the system

Use the **module_led_list** command to display the module LED state in the system.

```
module_led_list [ module=ModuleNumber ]
```

## Parameters

| Name | Description | Mandatory | Default |
|---|---|---|---|
| **module** | Limits the listing to a specific module. | N | All LEDs in all modules. |

### Example:

```
module_led_list module=1:module:3
```

### Output:

```
Module       LED              State   Color
------------ ---------------- ------- -------
1:Module:3   Battery          Off     na
1:Module:3   CPU 1            Off     na
1:Module:3   CPU 2            Off     na
1:Module:3   CPU Mismatch     Off     na
1:Module:3   Check Log        Off     na
1:Module:3   DIMM 1           Off     na
1:Module:3   DIMM 10          Off     na
1:Module:3   DIMM 11          Off     na
1:Module:3   DIMM 12          Off     na
1:Module:3   DIMM 13          Off     na
1:Module:3   DIMM 14          Off     na
1:Module:3   DIMM 15          Off     na
1:Module:3   DIMM 16          Off     na
1:Module:3   DIMM 17          Off     na
1:Module:3   DIMM 18          Off     na
1:Module:3   DIMM 19          Off     na
1:Module:3   DIMM 2           Off     na
1:Module:3   DIMM 20          Off     na
1:Module:3   DIMM 21          Off     na
1:Module:3   DIMM 22          Off     na
1:Module:3   DIMM 23          Off     na
1:Module:3   DIMM 24          Off     na
1:Module:3   DIMM 3           Off     na
1:Module:3   DIMM 4           Off     na
1:Module:3   DIMM 5           Off     na
1:Module:3   DIMM 6           Off     na
1:Module:3   DIMM 7           Off     na
1:Module:3   DIMM 8           Off     na
1:Module:3   DIMM 9           Off     na
1:Module:3   Fan 1            Off     na
1:Module:3   Fan 2            Off     na
1:Module:3   Fan 3            Off     na
1:Module:3   Fan 4            Off     na
1:Module:3   Fan 5            Off     na
1:Module:3   Fan 6            Off     na
1:Module:3   Fan Riser1       Off     na
1:Module:3   Fan Riser2       Off     na
1:Module:3   Fault            Off     na
1:Module:3   IMM2 Heartbeat   Blink   Green
1:Module:3   Identify         Off     na
1:Module:3   Internal RAID    Off     na
1:Module:3   PCI 1            Off     na
1:Module:3   PCI 2            Off     na
1:Module:3   PCI 3            Off     na
1:Module:3   PCI 4            Off     na
1:Module:3   PCI 5            Off     na
1:Module:3   PCI 6            Off     na
1:Module:3   PCI 7            Off     na
1:Module:3   PCI 8            Off     na
1:Module:3   Power            On      Green
1:Module:3   SysBrd Fault     Off     na
```

| Field ID | Field output | Default position |
|---|---|---|
| **module** | Module | 1 |
| **led** | LED | 2 |
| **state** | State | 3 |

| Field ID | Field output | Default position |
|---|---|---|
| **color** | Color | 4 |
| **reported** | Event Active | N/A |

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Allowed |
| Security administrator | Disallowed |
| Read-only users | Allowed |
| Technicians | Allowed |

## Listing data disk devices in the system

Use the **disk_list** command to lists special vault device statuses.

```
disk_list [ module=ModuleNumber  | disk=ComponentId ]
```

### Parameters

| Name | Description | Mandatory | Default |
|---|---|---|---|
| **module** | Limits the listing to a specific module. | N | All data disk devices in all modules. |
| **disk** | Data disk devices for which special statuses are to be listed. | N | All data disk devices. |

This command lists the statuses of the data disk devices, including:

- Component generic status
- Data disk device capacity
- Model
- Serial

**Note:**

The temperature values are indicated in Celsius.

**Example:**

```
disk_list
```

**Output:**

```
Component ID          Status   Currently Functioning   Capacity   Target Status   Vendor
--------------------- -------- ---------------------- ---------- --------------- --------
1:disk:14:1           OK       yes                    250GB                      LENOVO-X
1:disk:14:2           OK       yes                    250GB                      LENOVO-X
1:disk:3:1            OK       yes                    250GB                      LENOVO-X
1:disk:3:2            OK       yes                    250GB                      LENOVO-X
1:disk:5:1            OK       yes                    250GB                      LENOVO-X
1:disk:5:2            OK       yes                    250GB                      LENOVO-X
1:disk:6:1            OK       yes                    250GB                      LENOVO-X
1:disk:6:2            OK       yes                    250GB                      LENOVO-X

Cont.:
Model             Serial     Firmware   FRU       Temperature   Encryption State
----------------- ---------- ---------- --------- ------------- ------------------
HUSMR1625ASS20E   0PVGJTPA   P4C9       00NA685   22            Ready
HUSMR1625ASS20E   0PVGHN6A   P4C9       00NA685   22            Ready
HUSMR1625ASS20E   0PVJ2PEA   P4C9       00NA685   21            Ready
HUSMR1625ASS20E   0PVJ251A   P4C9       00NA685   21            Ready
HUSMR1625ASS20E   0PVJ99UA   P4C9       00NA685   20            Ready
HUSMR1625ASS20E   0PVJRB5A   P4C9       00NA685   21            Ready
HUSMR1625ASS20E   0PVJS39A   P4C9       00NA685   20            Ready
HUSMR1625ASS20E   0PVJ9RAA   P4C9       00NA685   21            Ready
```

| Field ID | Field output | Default position |
|---|---|---|
| **component_id** | Component ID | 1 |
| **status** | Status | 2 |
| **currently_functioning** | Currently Functioning | 3 |
| **capacity_in_bytes** | Capacity | N/A |
| **capacity** | Capacity | 4 |
| **target_status** | Target Status | 5 |
| **vendor** | Vendor | 6 |
| **original_vendor** | Original Vendor | N/A |
| **model** | Model | 7 |
| **original_model** | Original Model | N/A |
| **serial** | Serial | 8 |
| **original_serial** | Original Serial | N/A |
| **firmware** | Firmware | 9 |
| **original_firmware** | Original Firmware | N/A |
| **part_number** | FRU | 10 |
| **original_part_number** | Original FRU | N/A |
| **temperature** | Temperature | 11 |
| **encryption_state** | Encryption State | 12 |
| **hw_mon_node_id** | Hw Node Owner | N/A |
| **requires_service** | Requires Service | N/A |
| **service_reason** | Service Reason | N/A |
| **revision** | Revision | N/A |
| **drive_pn** | Drive P/N | N/A |
| **original_drive_pn** | Original Drive P/N | N/A |
| **fru_pn** | FRU P/N | N/A |
| **original_fru_pn** | Original FRU P/N | N/A |
| **desc.bgd_scan** | Background Scan | N/A |

| Field ID | Field output | Default position |
|---|---|---|
| **desc.disk_id** | Device ID | N/A |
| **desc.last_sample_serial** | Last Sample Serial | N/A |
| **desc.last_sample_time** | Last Sample Time | N/A |
| **desc.power_is_on** | Power On | N/A |
| **desc.power_on_hours** | Power On Hours | N/A |
| **desc.power_on_minutes** | Power On Minutes | N/A |
| **desc.last_time_pom_was_mod** | Last Time Power On Minutes Was Modified | N/A |
| **desc.read_fail** | Read Fail | N/A |
| **desc.smart_code** | SMART Code | N/A |
| **desc.smart_fail** | SMART Fail | N/A |
| **desc.temperature_status.reported_severity** | Reported Temperature Severity | N/A |
| **desc.temperature_status.reported_temperature** | Reported Temperature | N/A |
| **desc.temperature_status.temperature** | Device Temperature | N/A |
| **desc.sw_encryption_active** | Software-Based Encryption Active | N/A |

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Allowed |
| Security administrator | Allowed |
| Read-only users | Allowed |
| Technicians | Allowed |

## Listing service statuses

Use the **service_list** command to list all service specific statuses.

```
service_list [ service=ComponentId ]
```

### Parameters

| Name | Description | Mandatory | Default |
|---|---|---|---|
| **service** | The service to be listed. | N | All services |

This command lists the statuses that apply to services. The list includes the following information:

- Component generic status
- Service on/failed
- Comment (optional)

**Example:**

```
service_list
```

**Output:**

```
Component ID          Status   Currently Functioning   Target Status
--------------------- -------- ---------------------- ---------------
1:Data:10             OK       yes
1:Data:11             OK       yes
1:Data:7              OK       yes
1:Data:9              OK       yes
1:Data_Reduction:10   OK       yes
1:Data_Reduction:11   OK       yes
1:Data_Reduction:7    OK       yes
1:Data_Reduction:9    OK       yes
1:Interface:10        OK       yes
1:Interface:11        OK       yes
1:Interface:9         OK       yes
1:Remote:10           OK       yes
1:Remote:11           OK       yes
1:Remote:9            OK       yes
```

| Field ID | Field output | Default position |
|---|---|---|
| **component_id** | Component ID | 1 |
| **status** | Status | 2 |
| **currently_functioning** | Currently Functioning | 3 |
| **target_status** | Target Status | 4 |

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Allowed |
| Security administrator | Disallowed |
| Read-only users | Allowed |
| Technicians | Allowed |

## Listing system components that require service

Use the **component_service_required_list** command to list system components and their status.

```
component_service_required_list [ component=ComponentId ] [ filter=<ALL|FAILED|NOTOK> ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **component** | N/A | Lists only this component. | N | All components. |
| **filter** | Enumeration | Filters the list to show only failed or only non-OK components. | N | ALL |

The list can be filtered to show only a specific component, all failed components, or all components in a non-OK state.

For status and configuration of specific component types, refer to the **_list** commands for specific components, such as: **module_list** or **switch_list**. The output is a list of components with the following information for each component:

- Component identification
- Component general status
- Indication about whether the component is currently functioning

**Example:**

```
component_service_required_list
```

**Output:**

```
Component ID            Status   Currently Functioning   Requires Service
----------------------- -------- ----------------------- ------------------
1:IB_Module_Port:12:2   Failed   no                      COMPONENT_TEST
1:IB_Module_Port:13:2   Failed   no                      COMPONENT_TEST
1:IB_Module_Port:8:2    Failed   no                      COMPONENT_TEST
1:IB_Switch_Port:2:12   Failed   no                      COMPONENT_TEST
1:IB_Switch_Port:2:13   Failed   no                      COMPONENT_TEST
1:IB_Switch_Port:2:8    Failed   no                      COMPONENT_TEST

Cont.:

Service Reason
--------------------------
IB_SWITCH_PHY_PORT_NOT_UP
IB_SWITCH_PHY_PORT_NOT_UP
IB_SWITCH_PHY_PORT_NOT_UP
IB_SWITCH_PHY_PORT_NOT_UP
IB_SWITCH_PHY_PORT_NOT_UP
IB_SWITCH_PHY_PORT_NOT_UP
```

| Field ID | Field output | Default position |
|---|---|---|
| **component_id** | Component ID | 1 |
| **status** | Status | 2 |
| **currently_functioning** | Currently Functioning | 3 |
| **requires_service** | Requires Service | 4 |
| **service_reason** | Service Reason | 5 |

## Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Allowed |
| Security administrator | Disallowed |
| Read-only users | Allowed |
| Technicians | Allowed |

# Listing trace snapshot on a module.

Use the **traces_snapshot_list** command to list trace snapshots on a module.

```
traces_snapshot_list module=ModuleNumber
```

## Parameters

| Name | Description | Mandatory |
|---|---|---|
| **module** | Component ID of the module to query. | Y |

| Field ID | Field output | Default position |
|---|---|---|
| **snapshot** | Snapshot Directories | 1 |

**Example:**

```
traces_snapshot_list  module=1:Module:9
```

**Output:**

```
Snapshots Directories
--------------------------------
1__20120802_1653_20120802_1713
```

## Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Disallowed |
| Read-only users | Disallowed |
| Technicians | Allowed |

# Creating a trace snapshot

Use the **traces_snapshot** command to create a trace snapshot.

```
traces_snapshot [ snapshot_back_time=MINUTES ] [ snapshot_delay_time=MINUTES ]
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **snapshot_delay_time** | Integer | Max delay between the request and snapshot creation. | N | no. Uses configuration misc.internal.auto_snapshot_trace.last_snapshot_minutes_delay field. |
| **snapshot_back_time** | Integer | Time back from the request time to include in the snapshot. | N | no. Uses configuration misc.internal.auto_snapshot_trace.snapshot_back_time field. |

**Example:**

```
traces_snapshot snapshot_back_time=60 snapshot_delay_time=1
```

## Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Disallowed |
| Read-only users | Disallowed |
| Technicians | Allowed |

# 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. |
| **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 |

© Copyright IBM Corp. 2016, 2022

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **interval** | Positive integer | The length of time in each statistic's time point. The resolution of this number is set in *resolution_unit.* | Y | N/A |
| **resolution_unit** | Enumeration | Sets the unit of measurement for the length of each bin. | Y | N/A |
| **perf_class** | Object name | Displays performance class aggregated statistics for bandwidth and IOPS. | N | All Performance classes. |

This command lists I/O statistics. The **count** parameter sets the number of lines in the statistics report. The combination of the **interval** and **resolution_unit** parameters sets the length of time for each statistics line. Either start or end timestamp must be provided. These timestamps set the time for the statistics report. Other parameters restrict statistics to a specific host, host port, volume, domain, interface port and so on.

For each line of statistics, 48 numbers are reported, which represent all the combinations of reads/writes, hits/misses and I/O size reporting for each of the 16 options for bandwidth, IOPS and latency. Statistics collection is limited to 32 pools and 200 volumes.

The syntax for the **start** and **end** fields is as follows: `Y-M-D[.[h[:m[:s]]]]`, where the ranges are as follows:

- Y - year (four digit)
- M - month (1-12)
- D - day (1-31)
- h - hour (0-23, with 0 as default)
- m - minute (0-59, with 0 as default)
- s - second (0-59, with 0 as default)

The year, month and day are separated by dashes, and the optional hours, minutes and seconds are separated by colons.

Output units:

- Very Large blocks are >512KB
- Large blocks - 64-512KB
- Medium blocks - 8-64KB
- Small blocks - 0-8KB
- The latency is in Microseconds
- The bandwidth is in KB

| Field ID | Field output | Default position |
|---|---|---|
| **time** | Time | 1 |
| **failures** | Failures | N/A |
| **aborts** | Aborts | N/A |
| **read_hit_very_large_iops** | Read Hit Very large - IOps | 2 |
| **read_hit_very_large_latency** | Read Hit Very large - Latency | 3 |
| **read_hit_very_large_internal_latency** | Read Hit Very large - Internal Latency | 75 |
| **read_hit_very_large_throughput** | Read Hit Very large - Throughput | 4 |

| Field ID | Field output | Default position |
|---|---|---|
| **read_hit_very_large_remotely_served** | Read Hit Very large - Remotely Served IOs | 63 |
| **read_hit_large_iops** | Read Hit Large - IOps | 5 |
| **read_hit_large_latency** | Read Hit Large - Latency | 6 |
| **read_hit_large_internal_latency** | Read Hit Large - Internal Latency | 76 |
| **read_hit_large_throughput** | Read Hit Large - Throughput | 7 |
| **read_hit_large_remotely_served** | Read Hit Large - Remotely Served IOs | 64 |
| **read_hit_medium_iops** | Read Hit Medium - IOps | 8 |
| **read_hit_medium_latency** | Read Hit Medium - Latency | 9 |
| **read_hit_medium_internal_latency** | Read Hit Medium - Internal Latency | 77 |
| **read_hit_medium_throughput** | Read Hit Medium - Throughput | 10 |
| **read_hit_medium_remotely_served** | Read Hit Medium - Remotely Served IOs | 65 |
| **read_hit_small_iops** | Read Hit Small - IOps | 11 |
| **read_hit_small_latency** | Read Hit Small - Latency | 12 |
| **read_hit_small_internal_latency** | Read Hit Small - Internal Latency | 78 |
| **read_hit_small_throughput** | Read Hit Small - Throughput | 13 |
| **read_hit_small_remotely_served** | Read Hit Small - Remotely Served IOs | 66 |
| **read_miss_very_large_iops** | Read Miss Very large - IOps | 14 |
| **read_miss_very_large_latency** | Read Miss Very large - Latency | 15 |
| **read_miss_very_large_internal_latency** | Read Miss Very large - Internal Latency | 79 |
| **read_miss_very_large_throughput** | Read Miss Very large - Throughput | 16 |
| **read_miss_very_large_remotely_served** | Read Miss Very large - Remotely Served IOs | 67 |
| **read_miss_large_iops** | Read Miss Large - IOps | 17 |
| **read_miss_large_latency** | Read Miss Large - Latency | 18 |
| **read_miss_large_internal_latency** | Read Miss Large - Internal Latency | 80 |
| **read_miss_large_throughput** | Read Miss Large - Throughput | 19 |
| **read_miss_large_remotely_served** | Read Miss Large - Remotely Served IOs | 68 |
| **read_miss_medium_iops** | Read Miss Medium - IOps | 20 |
| **read_miss_medium_latency** | Read Miss Medium - Latency | 21 |
| **read_miss_medium_internal_latency** | Read Miss Medium - Internal Latency | 81 |
| **read_miss_medium_throughput** | Read Miss Medium - Throughput | 22 |
| **read_miss_medium_remotely_served** | Read Miss Medium - Remotely Served IOs | 69 |
| **read_miss_small_iops** | Read Miss Small - IOps | 23 |
| **read_miss_small_latency** | Read Miss Small - Latency | 24 |
| **read_miss_small_internal_latency** | Read Miss Small - Internal Latency | 82 |
| **read_miss_small_throughput** | Read Miss Small - Throughput | 25 |
| **read_miss_small_remotely_served** | Read Miss Small - Remotely Served IOs | 70 |

| Field ID | Field output | Default position |
|---|---|---|
| **write_hit_very_large_iops** | Write Hit Very large - IOps | 26 |
| **write_hit_very_large_latency** | Write Hit Very large - Latency | 27 |
| **write_hit_very_large_internal_latency** | Write Hit Very large - Internal Latency | 83 |
| **write_hit_very_large_throughput** | Write Hit Very large - Throughput | 28 |
| **write_hit_large_iops** | Write Hit Large - IOps | 29 |
| **write_hit_large_latency** | Write Hit Large - Latency | 30 |
| **write_hit_large_internal_latency** | Write Hit Large - Internal Latency | 84 |
| **write_hit_large_throughput** | Write Hit Large - Throughput | 31 |
| **write_hit_medium_iops** | Write Hit Medium - IOps | 32 |
| **write_hit_medium_latency** | Write Hit Medium - Latency | 33 |
| **write_hit_medium_internal_latency** | Write Hit Medium - Internal Latency | 85 |
| **write_hit_medium_throughput** | Write Hit Medium - Throughput | 34 |
| **write_hit_small_iops** | Write Hit Small - IOps | 35 |
| **write_hit_small_latency** | Write Hit Small - Latency | 36 |
| **write_hit_small_internal_latency** | Write Hit Small - Internal Latency | 86 |
| **write_hit_small_throughput** | Write Hit Small - Throughput | 37 |
| **write_miss_very_large_iops** | Write Miss Very large - IOps | 38 |
| **write_miss_very_large_latency** | Write Miss Very large - Latency | 39 |
| **write_miss_very_large_internal_latency** | Write Miss Very large - Internal Latency | 87 |
| **write_miss_very_large_throughput** | Write Miss Very large - Throughput | 40 |
| **write_miss_large_iops** | Write Miss Large - IOps | 41 |
| **write_miss_large_latency** | Write Miss Large - Latency | 42 |
| **write_miss_large_internal_latency** | Write Miss Large - Internal Latency | 88 |
| **write_miss_large_throughput** | Write Miss Large - Throughput | 43 |
| **write_miss_medium_iops** | Write Miss Medium - IOps | 44 |
| **write_miss_medium_latency** | Write Miss Medium - Latency | 45 |
| **write_miss_medium_internal_latency** | Write Miss Medium - Internal Latency | 89 |
| **write_miss_medium_throughput** | Write Miss Medium - Throughput | 46 |
| **write_miss_small_iops** | Write Miss Small - IOps | 47 |
| **write_miss_small_latency** | Write Miss Small - Latency | 48 |
| **write_miss_small_internal_latency** | Write Miss Small - Internal Latency | 90 |
| **write_miss_small_throughput** | Write Miss Small - Throughput | 49 |
| **read_memory_hit_very_large_iops** | Read Memory-Hit Very large - IOps | 50 |
| **read_memory_hit_very_large_latency** | Read Memory-Hit Very large - Latency | 51 |

| Field ID | Field output | Default position |
|---|---|---|
| **read_memory_hit_very_large_internal_latency** | Read Memory-Hit Very large - Internal Latency | 91 |
| **read_memory_hit_very_large_throughput** | Read Memory-Hit Very large - Throughput | 52 |
| **read_memory_hit_very_large_remotely_served** | Read Memory-Hit Very large - Remotely Served IOs | 71 |
| **read_memory_hit_large_iops** | Read Memory-Hit Large - IOps | 53 |
| **read_memory_hit_large_latency** | Read Memory-Hit Large - Latency | 54 |
| **read_memory_hit_large_internal_latency** | Read Memory-Hit Large - Internal Latency | 92 |
| **read_memory_hit_large_throughput** | Read Memory-Hit Large - Throughput | 55 |
| **read_memory_hit_large_remotely_served** | Read Memory-Hit Large - Remotely Served IOs | 72 |
| **read_memory_hit_medium_iops** | Read Memory-Hit Medium - IOps | 56 |
| **read_memory_hit_medium_latency** | Read Memory-Hit Medium - Latency | 57 |
| **read_memory_hit_medium_internal_latency** | Read Memory-Hit Medium - Internal Latency | 93 |
| **read_memory_hit_medium_throughput** | Read Memory-Hit Medium - Throughput | 58 |
| **read_memory_hit_medium_remotely_served** | Read Memory-Hit Medium - Remotely Served IOs | 73 |
| **read_memory_hit_small_iops** | Read Memory-Hit Small - IOps | 59 |
| **read_memory_hit_small_latency** | Read Memory-Hit Small - Latency | 60 |
| **read_memory_hit_small_internal_latency** | Read Memory-Hit Small - Internal Latency | 94 |
| **read_memory_hit_small_throughput** | Read Memory-Hit Small - Throughput | 61 |
| **read_memory_hit_small_remotely_served** | Read Memory-Hit Small - Remotely Served IOs | 74 |
| **time_in_seconds** | Time (s) | 62 |

## Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Allowed |
| Security administrator | Disallowed |
| Read-only users | Allowed |
| Technicians | Disallowed |

## Return codes

- **BAD_TIME_FORMAT**

  Bad time format. Should be YYYY-MM-DD[.HH[:MM[:SS]]].
- **TARGET_PORT_BAD_ADDRESS**

The remote port address is illegal or does not belong to the remote target.

- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **STATS_TOO_MANY_SAMPLES**

  The requested number of statistics samples is too high.
- **TARGET_BAD_NAME**

  The target name does not exist.
- **COMPONENT_DOES_NOT_EXIST**

  The component does not exist.
- **HOST_BAD_NAME**

  The host name does not exist.
- **HOST_PORT_DOES_NOT_EXIST**

  The port ID is not defined.
- **IPINTERFACE_DOES_NOT_EXIST**

  This IP interface name does not exist.
- **PERF_CLASS_BAD_NAME**

  The performance class does not exist.
- **COMMAND_AMBIGUOUS**

  The user belongs to more than one domain. Please specify a domain or an object.
- **DOMAIN_DOESNT_EXIST**

  The domain does not exist.

# Retrieving usage history

Use the **usage_get** command to display the usage history of a volume or a storage pool.

```
usage_get < vol=VolName | pool=PoolName > [ start=TimeStamp | start_in_seconds=StartTime ]
[ end=TimeStamp ]
 [ max=MaxEntries ]
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **vol** | Object name | Volume for which usage statistics are retrieved. | N | N/A |
| **pool** | Object name | Storage pool for which usage statistics are retrieved. | N | N/A |
| **start** | N/A | Starting time for usage history retrieval. | N | Creation time of the object. |
| **end** | N/A | Ending time for usage history retrieval. | N | Current time. |
| **max** | Integer | Maximum number of entries to retrieve. | N | No limit. |
| **start_in_seconds** | Integer | Starting time for usage history retrieval, in seconds since 12:00:00 AM, 1 January 1970. | N | Creation time of the object. |

This command retrieves the usage history of a storage pool or volume in megabytes (MB).

**Example:**

```
usage_get pool=DBPool
```

**Output:**

```
Time                 Volume Usage (MiB)   Snapshot Usage (MiB)
-------------------- -------------------- ----------------------
2016-03-29 12:00:00   0                    0
2016-03-29 13:00:00   0                    0
2016-03-29 14:00:00   0                    0
```

| Field ID | Field output | Default position |
|---|---|---|
| **time** | Time | 1 |
| **volume_usage** | Volume Usage (MiB) | 2 |
| **snapshot_usage** | Snapshot Usage (MiB) | 3 |

## Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Disallowed |
| Read-only users | Allowed |
| Technicians | Disallowed |

## Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **POOL_DOES_NOT_EXIST**

  The storage pool does not exist.
- **BAD_TIME_FORMAT**

  Bad time format. Should be YYYY-MM-DD[.HH[:MM[:SS]]].
- **END_BEFORE_START**

  The end time cannot precede the start time.
- **VOLUME_IS_SNAPSHOT**

  The operation is not permitted on snapshots.

# 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 |
| 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 |

© Copyright IBM Corp. 2016, 2022

### Return codes

- **OBJECT_BAD_NAME**

  The referenced object does not exist.
- **MAX_METADATA_OBJECTS_REACHED**

  The maximum number of metadata objects has been reached.
- **REMOTE_MAX_METADATA_OBJECTS_REACHED**

  The maximum number of metadata objects has been reached on a remote system.
- **LOCAL_PEER_IS_NOT_MASTER**

  The local peer is not primary.
- **TARGET_NOT_CONNECTED**

  There is currently no connection to the target system.
- **REMOTE_TARGET_NOT_CONNECTED**

  There is currently no connection from the target system.
- **HA_IS_NOT_OPERATIONAL**

  This HyperSwap relation is not operational. The operation cannot be carried out on a non-operational HyperSwap relation.
- **MAX_METADATA_VOL_OBJECTS_REACHED**

  The maximum number of metadata objects for a volume has been reached.

## Deleting metadata

Use the **metadata_delete** command to delete an object's metadata.

```
metadata_delete object_type=Object name=Name key=Key
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **object_type** | Enumeration | Type of object. Available values: cg, cluster, dest, destgroup, host, performanceclass, pool, rule, schedule, smsgw, smtpgw, target, user, user_group, vol. | Y |
| **name** | Object name | The name of the target object. | Y |
| **key** | String | Metadata key. | Y |

This command deletes a metadata key value for the specified object.

The command will fail if the key is not defined.

**Example:**

```
metadata_delete object_type=host name=Host1 key=01
```

**Output:**

```
Command completed successfully.
```

### Access control

| User Category | Permission | Condition |
|---|---|---|
| Storage administrator | Allowed | N/A |
| Storage integration administrator | Allowed | N/A |
| Application administrator | Conditionally Allowed | Metadata can be set for only volumes, snapshots, snapshot groups, clusters or hosts, and only for objects associated with the application administrator executing the command. Hosts or clusters should be associated with the user. Volumes should be mapped to a host or a cluster associated with the user. Snapshots or snapshot groups should be ones created by application administrator. |
| Security administrator | Disallowed | N/A |
| Read-only users | Disallowed | N/A |
| Technicians | Disallowed | N/A |

### Return codes

- **OBJECT_BAD_NAME**

  The referenced object does not exist.
- **METADATA_OBJECT_KEY_NOT_FOUND**

  The specified metadata object does not exist.
- **LOCAL_PEER_IS_NOT_MASTER**

  The local peer is not primary.
- **HA_IS_NOT_OPERATIONAL**

  This HyperSwap relation is not operational. The operation cannot be carried out on a non-operational HyperSwap relation.
- **TARGET_NOT_CONNECTED**

  There is currently no connection to the target system.
- **REMOTE_TARGET_NOT_CONNECTED**

  There is currently no connection from the target system.

## Listing metadata

Use the **metadata_list** command to list an object's metadata.

```
metadata_list [ object_type=Object ] [ name=Name ] [ key=Key ] [ domain=DomainName ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **object_type** | Enumeration | Type of object. | N | Type of object. Available values: cg, cluster, dest, destgroup, host, performanceclass, pool, rule, schedule, smsgw, smtpgw, target, user, user_group, vol. |

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **name** | Object name | The name of the target object. | N | All objects |
| **key** | String | Metadata key. | N | List all keys and values. |
| **domain** | Object name | The domain name. | N | All Domains |

This command lists all the value key pairs for this object, or a specific one. The command fails if no key is defined.

**Example:**

```
metadata_list object_type=host
```

**Output:**

```
Object Type   Name    Key   Value
------------- ------- ----- -------
Host          Host1   01    Host
```

| Field ID | Field output | Default position |
|---|---|---|
| **object_type** | Object Type | 1 |
| **name** | Name | 2 |
| **key** | Key | 3 |
| **value** | Value | 4 |

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Allowed |
| Application administrator | Allowed |
| Security administrator | Disallowed |
| Read-only users | Allowed |
| Technicians | Allowed |

## Setting user metadata

Use the **user_metadata_set** command to set user metadata by section and key.

```
user_metadata_set  section=Section key=Key value=Value
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **section** | Enumeration | Metadata section. | Y |
| **key** | String | Metadata key. | Y |
| **value** | String | Metadata value. | Y |

**Example:**

```
user_metadata_set section=GUI key=GUI_10.2 value=Ubunto_2.6
```

**Output:**

```
(Command returns no output)
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Disallowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

### Return codes

- **USER_METADATA_FULL**

  Maximal number of user metadata objects has been reached.

## Listing user metadata

Use the **user_metadata_list** command to list User meta data.

```
user_metadata_list [ section=Section ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **section** | Enumeration | User metadata section. | N | '' |

Lists User data.

**Example:**

```
user_metadata_list
```

**Output:**

```
Time                 Section  Key       Value
-------------------- -------- --------- -------
2004-11-22 18:08:23  TA       GUI_0.2   Ubunto_2.2
2007-11-22 18:08:22  GUI      GUI_10.1  Ubunto_2.6
2007-11-22 18:08:23  GUI      GUI_10.2  Ubunto_2.6
```

| Field ID | Field output | Default position |
|---|---|---|
| **time** | Time | 1 |
| **section** | Section | 2 |
| **key** | Key | 3 |

| Field ID | Field output | Default position |
|---|---|---|
| **value** | Value | 4 |

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Allowed |
| Security administrator | Disallowed |
| Read-only users | Allowed |
| Technicians | Disallowed |

## Deleting user metadata

Use the **user_metadata_delete** command to delete user metadata by user section and key.

```
user_metadata_delete section=Section key=Key
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **section** | Enumeration | Meta data section. | Y |
| **key** | String | Metadata key. | Y |

**Example:**

```
user_metadata_delete section=GUI key=GUI_10.2
```

**Output:**

```
(Command returns no output)
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Disallowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

### Return codes

- **USER_METADATA_OBJECT_KEY_NOT_FOUND**

  The specified user metadata object does not exist.

# 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 have failed
- No Flash enclosure is undergoing CCL

When data protection is disabled, a cryptographic erase is performed on all protected bands (ensuring that all existing user data is no longer accessible). After the command successfully completes, all bands are left in the unlocked state.

Disabling encryption when the encryption state is other than Active (displayed as `Enabled` in **state_list**) will result in an error.

**Example:**

```
encrypt_disable -y
```

**Output:**

```
Command completed successfully.
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Disallowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Allowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

### Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_DISABLE_ENCRYPTION**

  Are you sure you want to disable encryption on this system?

  **Troubleshooting:** To proceed with issuing the command, enter -y.

© Copyright IBM Corp. 2016, 2022

## Return codes

- **CANNOT_DISABLE_ENCRYPTION_WHILE_NOT_IN_FULLY_PROTECTED_STATE**

  Cannot disable encryption while not in a fully protected state.

  **Troubleshooting:** Resolve any issues preventing system from reaching full data protection state. Contact IBM Support.
- **UNSUPPORTED_HARDWARE**

  Cannot encrypt on unsupported hardware.

  **Troubleshooting:** Contact IBM support to verify encryption status.
- **ENCRYPT_NOT_ENABLED**

  Encryption is not enabled.

  **Troubleshooting:** Make sure that encryption is enabled and re-run the command.
- **VOLUME(S)_DEFINED**

  There are volumes defined, cannot disable encryption.

  **Troubleshooting:** All volumes must be removed before encryption is disabled.
- **CANNOT_UNMOUNT_STATISTIC_VOLUME**

  Failed to unmount the statistics volume for disabling encryption.

  **Troubleshooting:** Contact IBM Support.
- **DATA_REDUCTION_RECOVERY_IS_RUNNING**

  Data reduction recovery is running, the operation is not allowed.

  **Troubleshooting:** Wait for data reduction recovery to complete.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support
- **SECURE_ERASE_IS_RUNNING**

  The operation not allowed while secure erase is running.

  **Troubleshooting:** Wait for the secure erase process to complete.
- **FLASH_CCL_IN_PROGRESS**

  The requested command cannot be invoked while Flash Enclosure CCL is in progress.

  **Troubleshooting:** Wait for Flash CCL to complete.
- **ENCRYPTION_CANNOT_UNENROLL_SOME_VAULT_DEVICES**

  Some vault devices cannot be un-enrolled due to failed components.

  **Troubleshooting:** Contact IBM Support.
- **ENCRYPTION_CANNOT_UNENROLL_SOME_FLASH_ENCLOSURES**

  Not all flash enclosures are phased in before un-enrollment.

  **Troubleshooting:** Contact IBM Support.
- **ENCRYPTION_IS_NOT_IN_A_STABLE_STATE**

  Encryption is not in a stable state.

  **Troubleshooting:** Wait for the encryption process to complete.

# Enabling encryption

Use the **encrypt_enable** command to enable the data protection feature.

```
encrypt_enable [ recovery_keys=<yes|no> ] [ key_scheme=KeyScheme ]
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **recovery_keys** | Boolean | Defines whether recovery keys are required for encryption activation. | N | yes |
| **key_scheme** | Enumeration | Defines which key scheme to use for encryption activation: external or local. | N | external |

This command enables the data protection feature. In order for this command to complete successfully, all of the following prerequisites must be fulfilled:

- The system is fully redundant
- None of the present and active SSDs, Flash cards, or Flash canisters have failed
- No Flash enclosure is undergoing CCL

**Example:**

```
encrypt_enable recovery_keys=yes key_scheme=local -y
```

**Output:**

```
Command completed successfully.
```

## Access control

| User Category | Permission |
|---|---|
| Storage administrator | Disallowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Allowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

## Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_ENABLE_ENCRYPTION**

  Are you sure you want to enable encryption on this system?

  **Troubleshooting:** To proceed with issuing the command, enter -y.
- **ARE_YOU_SURE_YOU_WANT_TO_ENABLE_LOCAL_KEY_MANAGEMENT_ENCRYPTION**

  Are you sure you want to enable encryption with local key management on this system?

  **Troubleshooting:** To proceed with issuing the command, enter -y.

## Return codes

- **CANNOT_ENABLE_ENCRYPTION_WHILE_NOT_IN_FULLY_PROTECTED_STATE**

  Cannot enable encryption while not in a fully protected state.

  **Troubleshooting:** Resolve any issues preventing the system from reaching a full data protection state, and contact IBM Support.
- **UNSUPPORTED_HARDWARE**

  Cannot encrypt on unsupported hardware.

  **Troubleshooting:** Contact IBM support to verify encryption status.
- **INVALID_RECOVERY_KEY_STATE**

  The recovery key state is inconsistent with the specified option.

  **Troubleshooting:** Check the recovery key state using the encrypt_recovery_key_status command.
- **ENCRYPTION_ALREADY_ENABLED**

  Encryption has already been enabled.

  **Troubleshooting:** Check the output of the state_list command.
- **SYSTEM_IS_REDISTRIBUTING**

  The operation is not allowed during rebuild or phase-in.
- **NO_LIVE_KEYSERVER_GATEWAY_NODE**

  There is no live key server gateway node in the system.

  **Troubleshooting:** Restart the key server gateway node and try again.
- **NO_MASTER_KEYSERVER_DEFINED**

  No master key server is defined in the system.

  **Troubleshooting:** Define a master key server by invoking encrypt_key server_update and try again.
- **KEYSERVER_COMMUNICATION_GENERIC_ERROR**

  Cannot connect to an active key server.

  **Troubleshooting:** Invoke encrypt_keyserver_list and event_list for more details.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support
- **ENCRYPTION_CANNOT_ENROLL_SOME_FLASH_ENCLOSURES**

  Not all flash enclosures are phased in before enrollment.

  **Troubleshooting:** Contact IBM Support.
- **ENCRYPTION_CANNOT_ENROLL_SOME_VAULT_DEVICES**

  Some vault devices cannot be enrolled due to failed components.

  **Troubleshooting:** Contact IBM Support.
- **DATA_REDUCTION_RECOVERY_IS_RUNNING**

  Data reduction recovery is running, the operation is not allowed.

  **Troubleshooting:** Wait for data reduction recovery to complete.
- **FLASH_CCL_IN_PROGRESS**

  The requested command cannot be invoked while Flash Enclosure CCL is in progress.

  **Troubleshooting:** Wait for Flash CCL to complete.

- **RECOVERY_MISMATCH_KEY_SCHEME**

  The recovery key scheme does not match the current scheme.

  **Troubleshooting:** Check the recovery key scheme using the encrypt_key_scheme_get command.

# Defining a key server

Use the **encrypt_keyserver_define** command to define a key server to be used by the system.

```
encrypt_keyserver_define name=Name [ ipv4=Address ] [ ipv6=Address ] [ port=PortNumber ]
[ master=<yes|no> ] [ keyserver_type=KeyserverType ] certificate=PemCertificate
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **name** | String | The name of the key server being added. | Y | N/A |
| **certificate** | N/A | The public certificate or certificate chain of the key server being added (see below for details). | Y | N/A |
| **master** | Boolean | Defines whether this key server is the primary key server used for key retrieval. | N | no |
| **ipv4** | N/A | The IPv4 address of the key server being added. Either one IPv4 and/or one IPv6 must be used. | N | NONE |
| **ipv6** | N/A | The IPv6 address of the key server being added. Either one IPv4 and/or one IPv6 must be used. | N | NONE |
| **port** | Integer | Port used for key server communication. | N | 5696 |
| **keyserver_type** | Enumeration | The type of the key server to communicate with.<br>• Use **'TKLM'** for IBM TKLM, SKLM and iSKLM keyservers<br>• Use **'KeySecure'** for SafeNet, Gemalto KeySecure, Thales KeySecure, and Thales CipherTrust keyservers. | N | TKLM |

This command defines a key server to be used by the system upon startup or encryption activation to retrieve the key material required to cryptographically unlock the disks. At least one key server (but preferably two, and no more than four) must be defined and accessible in order for **encrypt_enable** to succeed. Only one of the key servers may be defined as master.

**The certificate parameter:**

The value of the **certificate** parameter is the content of a PEM file with asterisks instead of newlines. Chained certificates are supported. The total maximal length of a PEM file holding chained certificates (leaf first, root last) is 15360 characters (including the asterisk characters). In Windows, you can drag-and-drop a PEM file from the Windows Explorer to the appropriate location in the CLI session window; the content will be added automatically.

**Example:**

```
encrypt_keyserver_define name=snocone ipv4=snocone.example.com ipv6=2001:ODB8::a5a7 certificate=
"----BEGIN CERTIFICATE---
*MIICyTCCAbGgAwIBAgIXLSiyd2FPMA0GCSqGSIb3IiEBCwUAMBQxEjAQAgNVBVuTCXNrbG5pdHNv*
......
*erD5HgQHSkfR3FEM+b6EBOUPFIBrys8rKtLRbWvovobq*---END CERTIFICATE----"
```

**Output:**

```
Command completed successfully.
```

## Access control

| User Category | Permission |
|---|---|
| Storage administrator | Disallowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Allowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

## Return codes

- **UNSUPPORTED_HARDWARE**

  Cannot encrypt on unsupported hardware.

  **Troubleshooting:** Contact IBM support to verify encryption status.
- **ENCRYPTION_TOO_MANY_KEYSERVERS**

  Too many key servers are already defined, cannot add another one.

  **Troubleshooting:** Delete a key server and try again.
- **ENCRYPTION_KEYSERVER_NAME_EXISTS**

  The key server name already exists.

  **Troubleshooting:** Check the currently defined key servers.
- **ENCRYPTION_KEYSERVER_MUST_HAVE_ADDRESS**

  A key server must have at least one address (IPv4/IPv6).

  **Troubleshooting:** Make sure the command includes the ipv4= or ipv6= argument.
- **ENCRYPTION_KEYSERVER_IPV4_ALREADY_EXISTS**

  The IPv4 address or host name already exists.

  **Troubleshooting:** Check the currently defined key servers.
- **ENCRYPTION_KEYSERVER_IPV6_ALREADY_EXISTS**

  The IPv6 address or host name already exists.

  **Troubleshooting:** Check the currently defined key servers.
- **SSL_CERTIFICATE_HAS_EXPIRED**

  The SSL certificate has expired.
- **SSL_CERTIFICATE_VERIFICATION_FAILED**

  The SSL certificate chain verification failed.

- **SSL_CERTIFICATE_INVALID_FORMAT**

  The SSL certificate format is invalid or corrupted.
- **SSL_CERTIFICATE_NOT_YET_VALID**

  The SSL certificate is not yet valid.
- **SSL_CERTIFICATE_VERIFICATION_INTERNAL_ERROR**

  The SSL certificate verification has failed because of an internal system error.
- **SSL_CERTIFICATE_ISSUER_NOT_FOUND**

  The SSL certificate issuer was not found in the certificate chain.
- **SSL_CERTIFICATE_CHAIN_EMPTY**

  No certificates were found in the input.

## Removing a key server

Use the **encrypt_keyserver_delete** command to remove a key server used by the system.

```
encrypt_keyserver_delete name=Name
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **name** | String | The name of a defined key server. | Y |

**Example:**

```
encrypt_keyserver_delete name=snocone
```

**Output:**

```
Command completed successfully.
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Disallowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Allowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

### Return codes

- **UNSUPPORTED_HARDWARE**

  Cannot encrypt on unsupported hardware.

  **Troubleshooting:** Contact IBM support to verify encryption status.
- **ENCRYPTION_UNKNOWN_KEYSERVER**

  Unknown key server name.

**Troubleshooting:** Check the currently defined key servers.

- **ENCRYPTION_DELETE_MASTER_KEYSERVER**

  Removal of the master key server is not permitted.

  **Troubleshooting:** A new master key server must be defined before removing the current master.

- **ENCRYPTION_LAST_DEFINED_KEYSERVER**

  Cannot delete the last key server.

  **Troubleshooting:** Define another master key server before attempting to delete this one.

# Displaying key server status

Use the **encrypt_keyserver_list** command to list the key servers currently defined in the system along with their connectivity status.

```
encrypt_keyserver_list
```

**Example:**

```
encrypt_keyserver_list
```

**Output:**

```
Module   Name      App/Key Status   Last time checked   Master   Port
3        nachos    NOAPP            2013/03/27 20:18:43 yes      5696
3        nachos    UNKNOWN          2013/03/27 20:18:43 yes      5696
3        snocone   UNKNOWN          2013/03/27 20:18:43 no       5696
3        snocone   ACTIVE           2013/03/27 20:18:43 no       5696
3        TKLM-SA   BAD_CERT         2013/03/27 20:18:43 no       5696

Address
192.0.2.1
2001:0DB8::1
snocone.example.com
2001:0DB8:e006:238:209:6bff:fe00:a5a7
tklm-sa.example.com
```

| Field ID | Field output | Default position |
|---|---|---|
| **module_id** | Module | 1 |
| **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 |

| User Category | Permission |
|---|---|
| Read-only users | Allowed |
| Technicians | Allowed |

### Return codes

- **UNSUPPORTED_HARDWARE**

  Cannot encrypt on unsupported hardware.

  **Troubleshooting:** Contact IBM support to verify encryption status.

## Checking key server status

Use the **encrypt_keyserver_check_status** command to check connectivity status of the key servers currently defined in the system.

```
encrypt_keyserver_check_status
```

This command initiates the async check of the connectivity status of the key servers currently defined in the system. To get the current status, use the **encrypt_keyserver_list** CLI command.

**Example:**

```
encrypt_keyserver_check_status
```

**Output:**

```
Command completed successfully
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Allowed |
| Read-only users | Allowed |
| Technicians | Allowed |

### Return codes

- **UNSUPPORTED_HARDWARE**

  Cannot encrypt on unsupported hardware.

  **Troubleshooting:** Contact IBM support to verify encryption status.
- **CERTIFICATE_FOR_XIV_IS_NOT_INSTALLED**

  An XIV encryption certificate is not installed.

  **Troubleshooting:** Check the output of pki_list for a certificate named XIV and contact IBM Support.
- **ENCRYPTION_NOT_IN_EXTERNAL_SCHEME**

  Encryption key management is not set to the external scheme.

  **Troubleshooting:** Check the output of the encrypt_key_scheme_get command.

# Obtaining a new master key

Use the **encrypt_keyserver_rekey** command to initiate a rekey against the master key server.

```
encrypt_keyserver_rekey
```

This command initiates a rekeying (getting new cryptographic material) with the master key server.

**Example:**

```
encrypt_keyserver_rekey
```

**Output:**

```
Command completed successfully.
```

## Access control

| User Category | Permission |
|---|---|
| Storage administrator | Disallowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Allowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

## Return codes

- **UNSUPPORTED_HARDWARE**

  Cannot encrypt on unsupported hardware.

  **Troubleshooting:** Contact IBM support to verify encryption status.
- **ENCRYPT_NOT_ENABLED**

  Encryption is not enabled.

  **Troubleshooting:** Make sure that encryption is enabled and re-run the command.
- **CANNOT_GET_XIV_MASTER_KEY**

  Problem obtaining XIV master key from the key server.

  **Troubleshooting:** Make sure that the key service is active and serving keys (obtained from the key server or stored locally). Otherwise, contact IBM Support.
- **CANNOT_GET_NEW_KEY_REQUEST**

  Error requesting encryption keys from the key server gateway node.

  **Troubleshooting:** Make sure that the key server is actively serving keys.
- **CANNOT_UPDATE_KEY_METADATA**

  Cannot update metadata in the key repository for a new key.

  **Troubleshooting:** Contact IBM Support.
- **CANNOT_CANNOT_GENERATE_EXMK_ESKH**

  Problem generating EXMK and ESKH.

  **Troubleshooting:** Contact IBM Support.

- **CANNOT_WRITE_TO_KEY_REPOSITORY**

  Failed writing keys to the key repository.

  **Troubleshooting:** Contact IBM Support.
- **CANNOT_COPY_KEYS_IN_KEY_REPOSITORY**

  Problem copying current keys to the location of the old keys in the key repository.

  **Troubleshooting:** Contact IBM Support
- **ENCRYPTION_KR_WRITE_FAILED**

  Error writing to the key repository.

  **Troubleshooting:** Contact IBM Support.
- **ENCRYPTION_KR_READ_FAILED**

  Error reading the key repository.

  **Troubleshooting:** Contact IBM Support.
- **NO_LIVE_KEYSERVER_GATEWAY_NODE**

  There is no live key server gateway node in the system.

  **Troubleshooting:** Restart the key server gateway node and try again.
- **NO_MASTER_KEYSERVER_DEFINED**

  No master key server is defined in the system.

  **Troubleshooting:** Define a master key server by invoking encrypt_key server_update and try again.
- **KEYSERVER_COMMUNICATION_GENERIC_ERROR**

  Cannot connect to an active key server.

  **Troubleshooting:** Invoke encrypt_keyserver_list and event_list for more details.
- **ENCRYPTION_NOT_IN_EXTERNAL_SCHEME**

  Encryption key management is not set to the external scheme.

  **Troubleshooting:** Check the output of the encrypt_key_scheme_get command.

## Renaming a key server

Use the **encrypt_keyserver_rename** command to change the name of a defined key server.

```
encrypt_keyserver_rename name=Name new_name=Name
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **new_name** | String | The new name of the key server. | Y |
| **name** | String | The current name of a defined key server. | Y |

**Example:**

```
encrypt_keyserver_rename name=nachos new_name=snocone
```

**Output:**

```
Command completed successfully.
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Disallowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Allowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

### Return codes

- **UNSUPPORTED_HARDWARE**

  Cannot encrypt on unsupported hardware.

  **Troubleshooting:** Contact IBM support to verify encryption status.
- **ENCRYPTION_UNKNOWN_KEYSERVER**

  Unknown key server name.

  **Troubleshooting:** Check the currently defined key servers.
- **ENCRYPTION_KEYSERVER_NAME_EXISTS**

  The key server name already exists.

  **Troubleshooting:** Check the currently defined key servers.

## Changing key server properties

Use the **encrypt_keyserver_update** command to change a key server's IP address and/or port.

```
encrypt_keyserver_update name=Name [ ipv4=Address ] [ ipv6=Address ] [ port=PortNumber ]
[ master=<yes|no> ]
 [ certificate=PemCertificate ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **name** | String | Name of the key server to be updated. | Y | N/A |

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **certificate** | N/A | The public certificate or certificate chain of the key server to be updated.<br>The certificate parameter:<br>The value of the certificate parameter is the content of a PEM file with asterisks instead of newlines. Chained certificates are supported. The total maximal length of a PEM file holding chained certificates (leaf first, root last) is 15360 characters (including the asterisk characters). In Windows, you can drag-and-drop a PEM file from the Windows Explorer to the appropriate location in the XCLI session window; the content will be added automatically. | N | none |
| **master** | Enumeration | Indicates whether this key server is the master. | N | no |
| **ipv4** | N/A | The IPv4 address. | N | none |
| **ipv6** | N/A | The IPv6 address. | N | none |
| **port** | Integer | Port number for communications. | N | 5696 |

This command is used to update a key server's address, port, or certificate.

**Example:**

```
encrypt_keyserver_update name=nachos master=yes ipv4=192.0.2.1 ipv6=2001:ODB8::2 port=1010
certificate=''
```

**Output:**

```
Command completed successfully.
```

## Access control

| User Category | Permission |
|---|---|
| Storage administrator | Disallowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Allowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

### Return codes

- **UNSUPPORTED_HARDWARE**

  Cannot encrypt on unsupported hardware.

  **Troubleshooting:** Contact IBM support to verify encryption status.
- **ENCRYPTION_UNKNOWN_KEYSERVER**

  Unknown key server name.

  **Troubleshooting:** Check the currently defined key servers.
- **ENCRYPTION_KEYSERVER_IPV4_ALREADY_EXISTS**

  The IPv4 address or host name already exists.

  **Troubleshooting:** Check the currently defined key servers.
- **ENCRYPTION_KEYSERVER_IPV6_ALREADY_EXISTS**

  The IPv6 address or host name already exists.

  **Troubleshooting:** Check the currently defined key servers.
- **SSL_CERTIFICATE_HAS_EXPIRED**

  The SSL certificate has expired.
- **SSL_CERTIFICATE_VERIFICATION_FAILED**

  The SSL certificate chain verification failed.
- **SSL_CERTIFICATE_INVALID_FORMAT**

  The SSL certificate format is invalid or corrupted.
- **SSL_CERTIFICATE_NOT_YET_VALID**

  The SSL certificate is not yet valid.
- **SSL_CERTIFICATE_VERIFICATION_INTERNAL_ERROR**

  The SSL certificate verification has failed because of an internal system error.
- **SSL_CERTIFICATE_ISSUER_NOT_FOUND**

  The SSL certificate issuer was not found in the certificate chain.
- **SSL_CERTIFICATE_CHAIN_EMPTY**

  No certificates were found in the input.

## Entering a recovery key

Use the **encrypt_recovery_key_enter** command to unlock encrypted disks when the system reboots and cannot access any of the defined keyservers, and when recovery keys are defined.

```
encrypt_recovery_key_enter key=Key
```

### Parameters

| Name | Description | Mandatory |
|---|---|---|
| **key** | The 64-character hexadecimal recovery key. | Y |

This command is used to unlock encrypted disks when the system reboots and cannot access any of the defined keyservers. To unlock the disks, the **min_req number** (defined by the **encrypt_recovery_key_generate** command) of security administrators must all successfully enter their recovery key (as presented to them via **recovery_key_get**). After the minimum required keys

have been entered, the storage administrator must change the state from Maintenance to On by issuing **state_change target_state=on**. When this command is issued with the machine in the On state, it has no effect, and can be used to check the validity of the recovery key.

**Example:**

```
encrypt_recovery_key_enter key=CBC9B398373FDE79CD38B23192DABACADB5DA63A915CBF5CA8C4E0C212819DE6
```

**Output:**

```
Command completed successfully.
```

## Access control

| User Category | Permission |
|---|---|
| Storage administrator | Disallowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Allowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

## Return codes

- **UNSUPPORTED_HARDWARE**

  Cannot encrypt on unsupported hardware.

  **Troubleshooting:** Contact IBM support to verify encryption status.
- **INVALID_RECOVERY_KEY_FRAGMENT**

  The provided recovery key fragment does not match the stored key.

  **Troubleshooting:** Make sure that the proper key (share) has been used.
- **GENERIC_FAILED**

  Generic encryption failure.

  **Troubleshooting:** Contact IBM Support.
- **INVALID_RECOVERY_KEY_USER**

  The user is not a valid recovery key administrator.

  **Troubleshooting:** Make sure that the provided usernames are valid.
- **NO_LIVE_KEYSERVER_GATEWAY_NODE**

  There is no live key server gateway node in the system.

  **Troubleshooting:** Restart the key server gateway node and try again.
- **CANNOT_READ_FROM_KEY_REPOSITORY**

  Failed reading keys from the key repository.

  **Troubleshooting:** Contact IBM Support.
- **RK_FAILED_VERIFY_SLEEP**

  Too many failed verify attempts.

  **Troubleshooting:** Wait a little and try again.
- **ENCRYPTION_KR_WRITE_FAILED**

Error writing to the key repository.

**Troubleshooting:** Contact IBM Support.

- **RK_ENTER_SYSTEM_STATE_INVALID**

  The command is supported in maintenance mode only.

  **Troubleshooting:** Switch the system state to maintenance mode.

- **INVALID_RECOVERY_KEY_STATE**

  The recovery key state is inconsistent with the specified option.

  **Troubleshooting:** Check the recovery key state using the encrypt_recovery_key_status command.

- **RECOVERY_KEY_ALREADY_VERIFIED**

  The recovery key has already been verified.

  **Troubleshooting:** Check the recovery key state using the encrypt_recovery_key_list command.

## Generating recovery keys

Use the **encrypt_recovery_key_generate** command to specify which security administrators will receive recovery key shares, and to define the minimum number of recovery key shares that need to be entered.

```
encrypt_recovery_key_generate users=Users [ min_req=MinRequired ] [ key_scheme=KeyScheme ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **min_req** | Integer | Minimum number of required security administrator recovery key shares. | N | 2 |
| **users** | Object name | User names of the security administrators. | Y | N/A |
| **key_scheme** | Enumeration | Defines which key scheme to use for encryption activation external or local. | N | external |

This command is used to specify which security administrator will receive recovery keys (or, more accurately, "shares"), and to define the minimum number of recovery keys that need to be entered (using the **encrypt_recovery_key_enter** command) in order to unlock the encrypted keys. Once this command has been entered, all the specified security administrators are expected to retrieve and verify their recovery keys, using **encrypt_recovery_key_get** and **encrypt_recovery_key_verify**, respectively. This command can only be run when **encryption_state** is DISABLED.

**Example:**

```
encrypt_recovery_key_generate users=secadmin1,secadmin2,secadmin3,secadmin4 min_req=2
```

**Output:**

```
Command completed successfully.
```

## Access control

| User Category | Permission |
|---|---|
| Storage administrator | Disallowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Allowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

## Return codes

- **UNSUPPORTED_HARDWARE**

  Cannot encrypt on unsupported hardware.

  **Troubleshooting:** Contact IBM support to verify encryption status.
- **NO_LIVE_KEYSERVER_GATEWAY_NODE**

  There is no live key server gateway node in the system.

  **Troubleshooting:** Restart the key server gateway node and try again.
- **CANNOT_WRITE_TO_KEY_REPOSITORY**

  Failed writing keys to the key repository.

  **Troubleshooting:** Contact IBM Support.
- **CANNOT_GET_NEW_KEY_REQUEST**

  Error requesting encryption keys from the key server gateway node.

  **Troubleshooting:** Make sure that the key server is actively serving keys.
- **KEYSERVER_COMMUNICATION_GENERIC_ERROR**

  Cannot connect to an active key server.

  **Troubleshooting:** Invoke encrypt_keyserver_list and event_list for more details.
- **INSUFFICIENT_RK_ADMIN_THRESHOLD**

  Recovery key creation requires at least two security administrators.

  **Troubleshooting:** Re-run the command by indicating at least two security administrators.
- **ENCRYPTION_KR_WRITE_FAILED**

  Error writing to the key repository.

  **Troubleshooting:** Contact IBM Support.
- **ENCRYPTION_ALREADY_ENABLED**

  Encryption has already been enabled.

  **Troubleshooting:** Check the output of the state_list command.
- **NO_MASTER_KEYSERVER_DEFINED**

  No master key server is defined in the system.

  **Troubleshooting:** Define a master key server by invoking encrypt_key server_update and try again.
- **INVALID_RECOVERY_KEY_STATE**

  The recovery key state is inconsistent with the specified option.

  **Troubleshooting:** Check the recovery key state using the encrypt_recovery_key_status command.

- **INSUFFICIENT_RK_ADMINS**

  The number of users must be greater than or equal to the minimal required number.

  **Troubleshooting:** Re-run the command by providing at least the minimal number of required users.
- **CANNOT_GENERATE_KEYS_ON_KEYSERVER_GATEWAY**

  Failed to generate XMK and hashes on a key server gateway node.

  **Troubleshooting:** Contact IBM Support.
- **ENCRYPTION_KR_READ_FAILED**

  Error reading the key repository.

  **Troubleshooting:** Contact IBM Support.
- **CANNOT_UPDATE_KEY_METADATA**

  Cannot update metadata in the key repository for a new key.

  **Troubleshooting:** Contact IBM Support.

# Retrieving the security administrator's recovery key

Use the **encrypt_recovery_key_get** command to retrieve the recovery key share generated for the current user.

```
encrypt_recovery_key_get
```

This command retrieves the recovery key generated for the current user (by issuing **encrypt_recovery_key_generate** or **encrypt_recovery_key_rekey**) to be stored in a secure manner. After running this command, the user needs to 'prove' that they have the key by entering it via the **encrypt_recovery_key_verify** command. Once this is completed successfully, **encrypt_recovery_key_get** will no longer present the user's key. Using **encrypt_recovery_key_get** more than once will return the same value again.

**Example:**

```
encrypt_recovery_key_get
```

**Output:**

```
Command completed successfully.
key=B07C4374AC26C4DD3EC2E755EB3FAAF04EC792C8BE0D0CB1C1BAC79998EBEC6D
```

## Access control

| User Category | Permission |
|---|---|
| Storage administrator | Disallowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Allowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

## Return codes

- **UNSUPPORTED_HARDWARE**

Cannot encrypt on unsupported hardware.

**Troubleshooting:** Contact IBM support to verify encryption status.

- **INVALID_RECOVERY_KEY_USER**

  The user is not a valid recovery key administrator.

  **Troubleshooting:** Make sure that the provided usernames are valid.
- **NO_LIVE_KEYSERVER_GATEWAY_NODE**

  There is no live key server gateway node in the system.

  **Troubleshooting:** Restart the key server gateway node and try again.
- **CANNOT_READ_FROM_KEY_REPOSITORY**

  Failed reading keys from the key repository.

  **Troubleshooting:** Contact IBM Support.
- **CANNOT_GET_NEW_KEY_REQUEST**

  Error requesting encryption keys from the key server gateway node.

  **Troubleshooting:** Make sure that the key server is actively serving keys.
- **KEYSERVER_COMMUNICATION_GENERIC_ERROR**

  Cannot connect to an active key server.

  **Troubleshooting:** Invoke encrypt_keyserver_list and event_list for more details.
- **NO_MASTER_KEYSERVER_DEFINED**

  No master key server is defined in the system.

  **Troubleshooting:** Define a master key server by invoking encrypt_key server_update and try again.
- **INVALID_RECOVERY_KEY_STATE**

  The recovery key state is inconsistent with the specified option.

  **Troubleshooting:** Check the recovery key state using the encrypt_recovery_key_status command.
- **RECOVERY_KEY_ALREADY_VERIFIED**

  The recovery key has already been verified.

  **Troubleshooting:** Check the recovery key state using the encrypt_recovery_key_list command.

## Rekeying the security administrators

Use the **encrypt_recovery_key_rekey** command to restart the recovery key generation process.

```
encrypt_recovery_key_rekey [ users=Users ] [ min_req=MinRequired ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **min_req** | Integer | Minimum number of required security administrator recovery key shares. | N | 0 |
| **users** | Object name | Comma delimited list of security administrator to rekey. | N | N/A |

This command restarts the recovery key generation process, described in the section on the **encrypt_recovery_key_generate** command. The only difference is that the parameters **users**

and **min_required** are optional, and will default to the values specified in the last call to **encrypt_recovery_key_generate**. Note that none of the new recovery keys will take effect until the last user has verified his or her recovery key. Until then, if recovery is required, the previous keys will remain valid.

**Example:**

```
encrypt_recovery_key_rekey users=secadmin1,secadmin2,secadmin3,secadmin4 min_req=3
```

**Output:**

```
Command completed successfully.
```

## Access control

| User Category | Permission |
|---|---|
| Storage administrator | Disallowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Allowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

## Return codes

- **UNSUPPORTED_HARDWARE**

  Cannot encrypt on unsupported hardware.

  **Troubleshooting:** Contact IBM support to verify encryption status.
- **NO_LIVE_KEYSERVER_GATEWAY_NODE**

  There is no live key server gateway node in the system.

  **Troubleshooting:** Restart the key server gateway node and try again.
- **CANNOT_READ_FROM_KEY_REPOSITORY**

  Failed reading keys from the key repository.

  **Troubleshooting:** Contact IBM Support.
- **INSUFFICIENT_RK_ADMIN_THRESHOLD**

  Recovery key creation requires at least two security administrators.

  **Troubleshooting:** Re-run the command by indicating at least two security administrators.
- **ENCRYPTION_KR_WRITE_FAILED**

  Error writing to the key repository.

  **Troubleshooting:** Contact IBM Support.
- **NO_MASTER_KEYSERVER_DEFINED**

  No master key server is defined in the system.

  **Troubleshooting:** Define a master key server by invoking encrypt_key server_update and try again.
- **INVALID_RECOVERY_KEY_STATE**

  The recovery key state is inconsistent with the specified option.

  **Troubleshooting:** Check the recovery key state using the encrypt_recovery_key_status command.

- **INSUFFICIENT_RK_ADMINS**

  The number of users must be greater than or equal to the minimal required number.

  **Troubleshooting:** Re-run the command by providing at least the minimal number of required users.
- **CANNOT_GENERATE_KEYS_ON_KEYSERVER_GATEWAY**

  Failed to generate XMK and hashes on a key server gateway node.

  **Troubleshooting:** Contact IBM Support.
- **KEYSERVER_COMMUNICATION_GENERIC_ERROR**

  Cannot connect to an active key server.

  **Troubleshooting:** Invoke encrypt_keyserver_list and event_list for more details.

# Displaying recovery key status

Use the **encrypt_recovery_key_status** command to display status information for recovery keys.

```
encrypt_recovery_key_status
```

This command shows status information regarding recovery keys, specifically: Which user has verified his or her recovery key before **encrypt_enable** or in the recovery key rekey process. When using the recovery key to unlock the disks, which user has entered his or her recovery key. For information about the number of shares defined and the minimum number required for recovery, issue the **encrypt_recovery_key_list** command.

**Example:**

```
encrypt_recovery_key_status
```

**Output:**

```
Mon Aug 12 20:04:43 IDT 2013
Date Created          User        Status
2013-01-03 18:54:46   secadmin1   Verified
2013-01-03 18:54:46   secadmin2   Verified
2013-01-03 18:54:46   secadmin3   Verified
2013-01-03 18:54:46   secadmin4   Verified
2013-01-03 19:00:03   secadmin1   Unverified
2013-01-03 19:00:03   secadmin2   Unverified
2013-01-03 19:00:03   secadmin3   Unverified
2013-01-03 19:00:03   secadmin4   Unverified

When entering keys to unlock the disks:
Date Created          User        Status
2013-01-03 19:00:03   secadmin1   Accepted
2013-01-03 19:00:03   secadmin2   Accepted
2013-01-03 19:00:03   secadmin3   Pending
2013-01-03 19:00:03   secadmin4   Pending
```

| Field ID | Field output | Default position |
|---|---|---|
| **create_date** | Date Created | 1 |
| **user** | User | 2 |
| **status** | Status | 3 |

## Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |

| User Category | Permission |
|---|---|
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Allowed |
| Read-only users | Allowed |
| Technicians | Allowed |

### Return codes

- **UNSUPPORTED_HARDWARE**

  Cannot encrypt on unsupported hardware.

  **Troubleshooting:** Contact IBM support to verify encryption status.
- **CANNOT_READ_FROM_KEY_REPOSITORY**

  Failed reading keys from the key repository.

  **Troubleshooting:** Contact IBM Support.

## Recovering key verification

Use the **encrypt_recovery_key_verify** command to confirm that the current user has correctly copied the recovery key share retrieved by the **encrypt_recovery_key_get** command.

```
encrypt_recovery_key_verify key=Key
```

### Parameters

| Name | Description | Mandatory |
|---|---|---|
| **key** | The 64 character hexadecimal recovery key. | Y |

This command is used by security administrators to confirm that they have correctly copied the recovery key presented by the **encrypt_recovery_key_get** command. Encryption can be enabled (or a rekey can be completed) only when all security administrators have confirmed their respective recovery keys using this command.

**Example:**

```
encrypt_recovery_key_verify key=B07C4374AC26C4DD3EC2E755EB3FAAF04EC792C8BE0D0CB1C1BAC79998EBEC6D
```

**Output:**

```
Command completed successfully.
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Disallowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Allowed |
| Read-only users | Disallowed |

| User Category | Permission |
|---|---|
| Technicians | Disallowed |

### Return codes

- **UNSUPPORTED_HARDWARE**

  Cannot encrypt on unsupported hardware.

  **Troubleshooting:** Contact IBM support to verify encryption status.
- **INVALID_RECOVERY_KEY_FRAGMENT**

  The provided recovery key fragment does not match the stored key.

  **Troubleshooting:** Make sure that the proper key (share) has been used.
- **GENERIC_FAILED**

  Generic encryption failure.

  **Troubleshooting:** Contact IBM Support.
- **INVALID_RECOVERY_KEY_USER**

  The user is not a valid recovery key administrator.

  **Troubleshooting:** Make sure that the provided usernames are valid.
- **CANNOT_READ_FROM_KEY_REPOSITORY**

  Failed reading keys from the key repository.

  **Troubleshooting:** Contact IBM Support.
- **RK_FAILED_VERIFY_SLEEP**

  Too many failed verify attempts.

  **Troubleshooting:** Wait a little and try again.
- **ENCRYPTION_KR_WRITE_FAILED**

  Error writing to the key repository.

  **Troubleshooting:** Contact IBM Support.
- **INVALID_RECOVERY_KEY_STATE**

  The recovery key state is inconsistent with the specified option.

  **Troubleshooting:** Check the recovery key state using the encrypt_recovery_key_status command.
- **RECOVERY_KEY_ALREADY_VERIFIED**

  The recovery key has already been verified.

  **Troubleshooting:** Check the recovery key state using the encrypt_recovery_key_list command.

## Recovering key share information

Use the **encrypt_recovery_key_list** command to list recovery key share information.

```
encrypt_recovery_key_list
```

This command lists information regarding recovery keys, specifically: How many parts was the recovery key shared across, and how many are needed for the recovery process. When the currently valid recovery keys were created. To retrieve per-user information about the status of each key share, use the **encrypt_recovery_key_status** command.

**Example:**

```
encrypt_recovery_key_list
```

**Output:**

```
Recovery Key Initial Generation:

Date created        Number of Shares    Min Required
------------        ----------------    ------------
2013-03-11 16:00    3                   2

Recovery Key Rekeyed:
Date created        Number of Shares    Min Required
------------        ----------------    ------------
2013-03-11 16:00    3                   2
2013-03-20 16:05    4                   2
```

| Field ID | Field output | Default position |
|---|---|---|
| **create_date** | Key Created | 1 |
| **number_of_shares** | Number of Shares | 2 |
| **min_req** | Min Required | 3 |

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Allowed |
| Read-only users | Allowed |
| Technicians | Allowed |

### Return codes

- **UNSUPPORTED_HARDWARE**

  Cannot encrypt on unsupported hardware.

  **Troubleshooting:** Contact IBM support to verify encryption status.
- **CANNOT_READ_FROM_KEY_REPOSITORY**

  Failed reading keys from the key repository.

  **Troubleshooting:** Contact IBM Support.

## Finishing the recovery process

Use the **encrypt_recovery_finish** command to finish the recovery process and move the system to the On state.

```
encrypt_recovery_finish
```

Upon entering the recovery keys (see Entering a recovery key), this command finishes the recovery process and moves the system to the On state, provided that no more issues exist.

**Example:**

```
encrypt_recovery_finish
```

**Output:**

```
Command completed successfully.
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Allowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

### Return codes

- **ENCRYPT_NOT_ENABLED**

  Encryption is not enabled.

  **Troubleshooting:** Make sure that encryption is enabled and re-run the command.
- **RK_ENTER_SYSTEM_STATE_INVALID**

  The command is supported in maintenance mode only.

  **Troubleshooting:** Switch the system state to maintenance mode.

## Obtaining a new master key

Use the **encrypt_local_rekey** command to initiate rekeying (getting new cryptographic material) from the local key management.

```
encrypt_local_rekey
```

**Example:**

```
encrypt_local_rekey
```

**Output:**

```
Command completed successfully.
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Disallowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Allowed |
| Read-only users | Disallowed |

| User Category | Permission |
|---|---|
| Technicians | Disallowed |

## Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_REKEY**

  Are you sure you want to change the local key?

  **Troubleshooting:** To proceed with issuing the command, enter -y.

## Return codes

- **ENCRYPT_NOT_ENABLED**

  Encryption is not enabled.

  **Troubleshooting:** Make sure that encryption is enabled and re-run the command.
- **CANNOT_GET_XIV_MASTER_KEY**

  Problem obtaining XIV master key from the key server.

  **Troubleshooting:** Make sure that the key service is active and serving keys (obtained from the key server or stored locally). Otherwise, contact IBM Support.
- **CANNOT_GET_NEW_KEY_REQUEST**

  Error requesting encryption keys from the key server gateway node.

  **Troubleshooting:** Make sure that the key server is actively serving keys.
- **CANNOT_UPDATE_KEY_METADATA**

  Cannot update metadata in the key repository for a new key.

  **Troubleshooting:** Contact IBM Support.
- **CANNOT_CANNOT_GENERATE_EXMK_ESKH**

  Problem generating EXMK and ESKH.

  **Troubleshooting:** Contact IBM Support.
- **CANNOT_WRITE_TO_KEY_REPOSITORY**

  Failed writing keys to the key repository.

  **Troubleshooting:** Contact IBM Support.
- **CANNOT_COPY_KEYS_IN_KEY_REPOSITORY**

  Problem copying current keys to the location of the old keys in the key repository.

  **Troubleshooting:** Contact IBM Support
- **ENCRYPTION_KR_WRITE_FAILED**

  Error writing to the key repository.

  **Troubleshooting:** Contact IBM Support.
- **ENCRYPTION_KR_READ_FAILED**

  Error reading the key repository.

  **Troubleshooting:** Contact IBM Support.
- **NO_LIVE_KEYSERVER_GATEWAY_NODE**

  There is no live key server gateway node in the system.

  **Troubleshooting:** Restart the key server gateway node and try again.
- **ENCRYPTION_NOT_IN_LOCAL_SCHEME**

Encryption key management is not set to a local scheme.

**Troubleshooting:** Check the output of the encrypt_key_scheme_get command.

## Changing the key management scheme

Use the **encrypt_change_key_scheme** command to change the key management scheme.

```
encrypt_change_key_scheme key_scheme=KeyScheme
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **key_scheme** | Enumeration | Defines which key management scheme (external or local) to use for encryption activation. | Y |

This command only supports the change from an external to a local scheme.

**Example:**

```
encrypt_change_key_scheme key_scheme=local
```

**Output:**

```
Command completed successfully.
```

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Disallowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Allowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

### Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_CHANGE_KEY_SCHEME**

  Are you sure you want to change the key management scheme? This operation is irreversible.

  **Troubleshooting:** To proceed with issuing the command, enter -y.

### Return codes

- **ENCRYPT_NOT_ENABLED**

  Encryption is not enabled.

  **Troubleshooting:** Make sure that encryption is enabled and re-run the command.
- **CANNOT_GET_XIV_MASTER_KEY**

  Problem obtaining XIV master key from the key server.

**Troubleshooting:** Make sure that the key service is active and serving keys (obtained from the key server or stored locally). Otherwise, contact IBM Support.

- **CANNOT_GET_NEW_KEY_REQUEST**

  Error requesting encryption keys from the key server gateway node.

  **Troubleshooting:** Make sure that the key server is actively serving keys.
- **CANNOT_UPDATE_KEY_METADATA**

  Cannot update metadata in the key repository for a new key.

  **Troubleshooting:** Contact IBM Support.
- **CANNOT_CANNOT_GENERATE_EXMK_ESKH**

  Problem generating EXMK and ESKH.

  **Troubleshooting:** Contact IBM Support.
- **CANNOT_WRITE_TO_KEY_REPOSITORY**

  Failed writing keys to the key repository.

  **Troubleshooting:** Contact IBM Support.
- **CANNOT_COPY_KEYS_IN_KEY_REPOSITORY**

  Problem copying current keys to the location of the old keys in the key repository.

  **Troubleshooting:** Contact IBM Support
- **ENCRYPTION_KR_WRITE_FAILED**

  Error writing to the key repository.

  **Troubleshooting:** Contact IBM Support.
- **ENCRYPTION_KR_READ_FAILED**

  Error reading the key repository.

  **Troubleshooting:** Contact IBM Support.
- **NO_LIVE_KEYSERVER_GATEWAY_NODE**

  There is no live key server gateway node in the system.

  **Troubleshooting:** Restart the key server gateway node and try again.
- **ENCRYPTION_NOT_IN_EXTERNAL_SCHEME**

  Encryption key management is not set to the external scheme.

  **Troubleshooting:** Check the output of the encrypt_key_scheme_get command.

## Viewing the key scheme

Use the **encrypt_key_scheme_get** command to view the key scheme defined in the system.

```
encrypt_key_scheme_get
```

**Example:**

```
encrypt_key_scheme_get
```

**Output:**

```
Command completed successfully.
encrypt_key_scheme = "LOCAL"
```

## Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Allowed |
| Read-only users | Allowed |
| Technicians | Allowed |

# Chapter 25. Multi-site high availability and disaster recovery commands

This section describes the command-line interface (CLI) for Multi-site high availability and disaster recovery.

**Important:** The input and output syntax of CLI commands uses the legacy terminology of "Master", "SMaster", and "Slave" volumes, which in any documentation except the CLI reference, are referred to as "Primary", "Secondary", and "Tertiary". This inconsistency is a necessary compromise, required to avoid changes to older CLI commands that are in customer use, and also to keep the CLI terminology consistent across the board. The new terminology helps emphasize the commonality between the more recent functions of Multi-site HA/DR, high availability (HyperSwap), and the disaster recovery (Synchronous and Asynchronous mirroring) ones. It is used outside the CLI reference, where broader concepts can be explained.

## Viewing the Multi-site relation status

Use the **multisite_list** command to list the configuration and status of the defined Multi-site relations.

```
multisite_list [ vol=VolName | cg=cgName | scope=<cg|volume> ] [ domain=DomainName ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **scope** | Enumeration | All Multi-site relations, volume level Multi-site relations, CG level Multi-site relations. | N | All (if no value is specified) |
| **vol** | Object name | Local volume name. | N | [none] |
| **cg** | Object name | Local consistency group name. | N | [none] |
| **domain** | Object name | The domain name. | N | All Domains. |

If no parameters are specified, all the existing Multi-site relations are listed. When a name is indicated, only the specified Multi-site relation is listed.

The following parameters are listed by default:

- Name
- Multisite Object (Volume or CG)
- Multisite ID
- Role (Master, SMaster, Slave)
- State: This parameter can be retrieved only on the Master. The existing options are explained in the table below:

| Value | Meaning |
|---|---|
| `Inactive` | Both two-way relations that make part of a Multi-site relation are inactive. |
| `Initializing` | One or both two-way relations that make part of a Multi-site relation are in the init state. |

© Copyright IBM Corp. 2016, 2022

| Value | Meaning |
|---|---|
| Operational | The HyperSwap relation is synchronized, automatic failover is active, the asynchronous mirror is RPO OK, the SMaster is ready to assume the ownership of the disaster recovery (DR). |
| Degraded | The asynchronous mirror is RPO lagging. |
| Compromised | This state can be triggered by one of the following:<br>– The connection is disrupted between Master and SMaster and/or Master and Slave, or between SMaster and Slave. The latter is applicable only if a Standby mirror is configured.<br>– Resynchronization between Master and SMaster and/or Master and Slave is in progress.<br>– The machine holding the SMaster cannot assume the ownership of disaster recovery (DR), because the SMaster does not have an updated MRS (most recent snapshot) or LRS (last replicated snapshot.)<br>– A partial role change, either automatic or manual, has occurred between Master and SMaster or Master and Slave.<br>– The Slave cannot perform automatic failover. This may occur because: the Slave is not synchronized, the Slave is disconnected from the Quorum Witness, automatic failover is disabled by the user, or connectivity between Master and Slave has failed. |
| Inconsistent | A relation between Master and Slave is not defined. This can happen when a Standby relation was not created and the SMaster becomes the Master. |

- Multisite Standby: The state of the registered Standby mirror between SMaster and Slave. An asynchronous mirroring relation is considered a registered Standby mirror if it is defined on a Multi-site SMaster volume or consistency group, and is inactive and had either been registered or active in the past. The existing options are explained in the table below:

| Value | Meaning |
|---|---|
| Up | The Standby mirror is defined, registered and connected. |
| Down | The Standby mirror is defined and registered, but disconnected. |
| N/A | The Standby mirror is either undefined or unregistered. |

- **Master**: The name of the system that holds the Master volume or consistency group. For the local system, local is returned.
- **SMaster**: The name of the system that holds the SMaster volume or consistency group. For the local system, local is returned.
- **Slave**: The name of the system that holds the Slave volume or consistency group. For the local system, local is returned.
- **Designation**: The designated role of the local volume or consistency group in the Multi-site relation.

The following optional parameters can be listed by explicitly specifying the proper columns:

- **Main relation active**: Indicates whether the two-way relation between the Master and SMaster is active. The available options are Yes, No, or N/A. The latter will appear if the information is returned by the Slave.
- **Main relation sync state**: Indicates whether the Master and SMaster are synchronized. N/A will appear if the information is returned by the Slave.
- **Automatic failover**: Indicates whether automatic failover between Master and SMaster is available. Can be retrieved from either Master or SMaster. N/A will appear if the information is returned by the Slave.

- **Async mirror active**: Indicates whether the asynchronous mirroring between the Master and Slave is active. N/A will appear if the information is returned by the SMaster.
- **HA I/O Service**: Indicates the local peer's ability to serve I/O. The existing options are Active or Unavailable.
- **Standby snapshot state**: Indicates whether the SMaster has the data required for the Standby mirror. N/A will appear if the information is returned by the Slave.

**Example:**

```
multisite_list
```

**Output:**

```
Name   Multisite Object   Multisite ID       Role     State         Multisite Standby
Master          SMaster   Slave           Designation
------ ------------------ ------------------ -------- ------------- ----------------
-------------- --------- -------------- -------------
Va_1   Volume             5BC3473000000000   Master   Operational   Up
Local           Sync      DR              Master
Vb_2   Volume             5BC3473000000001   SMaster  N/A           Up
Sync            Local     DR              SMaster
Vc_2   Volume             5BC3473000000002   Slave    N/A           Up
DR              DR        Local           Slave
cg_a   CG                 5BC3473000000003   Master   Operational   Up
Local           Sync      DR              Master
```

| Field ID | Field output | Default position |
|---|---|---|
| **local_peer_name** | Name | 1 |
| **multisite_object** | Multisite Object | 2 |
| **multisite_global_id** | Multisite ID | 3 |
| **multisite_role** | Role | 4 |
| **multisite_state** | State | 5 |
| **standby_snapshot_state** | SMaster Mirror Ready | N/A |
| **standby_multisite_state** | Multisite Standby | 6 |
| **master_target_name** | Master | 7 |
| **smaster_target_name** | SMaster | 8 |
| **slave_target_name** | Slave | 9 |
| **designation** | Designation | 10 |
| **master_peer_name** | Master Peer | N/A |
| **smaster_peer_name** | SMaster Peer | N/A |
| **slave_peer_name** | Slave Peer | N/A |
| **main_relation_active** | Main relation active | N/A |
| **main_relation_sync_state** | Main relation sync state | N/A |
| **automatic_failover** | Automatic Failover | N/A |
| **async_mirror_active** | Async mirror active | N/A |
| **async_mirror_sync_state** | Async mirror sync state | N/A |
| **io_service** | HA I/O Service | N/A |

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Allowed |
| Application administrator | Allowed |
| Security administrator | Disallowed |
| Read-only users | Allowed |
| Technicians | Disallowed |

## Creating a Multi-site relation

Use the **multisite_define** command to create a new Multi-site relation.

```
multisite_define < vol=VolName | cg=cgName >
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **vol** | Object name | A volume name. | N |
| **cg** | Object name | Master consistency group name. | N |

Provided that a HyperSwap relation is established between Master and SMaster, and an asynchronous mirror is defined between Master and Slave, issuing this command on the Master creates a named Multi-site relation for a volume or consistency group. If the command completes successfully, the Multi-site relation will be recognized on all the systems involved.

The name of the Multi-site relation on all peers will be identical with the name of the target volume or consistency group on the Master. If a Multi-site relation with this name already exists on any system involved, the command will fail. The command will also fail if any of the following preliminary requirements is not met:

- All three systems are up and connected with the target.
- The target connectivity is configured as follows: Master-SMaster, SMaster-Master, Master-Slave. If the Standby relation is included, then SMaster-Slave connectivity must be configured and up.
- A HyperSwap relation is defined between Master and SMaster and vice versa. The HyperSwap relation status is either Active and Sync, or Inactive.
- An asynchronous mirroring relation is defined between Master and Slave and vice versa. The asynchronous mirroring relation status is either Active and Sync, or Inactive.
- Only one of the relations - either HyperSwap or asynchronous mirroring - is active.
- Either the HyperSwap or the asynchronous mirroring relation is defined as part of Multi-site.

Issuing **multisite_define** on a consistency group, automatically adds the Multi-site relation to each member volume.

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Allowed |
| Application administrator | Disallowed |

| User Category | Permission |
|---|---|
| Security administrator | Disallowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

## Warnings

- **DESIGNATION_MISMATCH_ARE_YOU_SURE_YOU_WANT_TO_OVERWRITE_EXISTING_DESIGNATION**

  There is a mismatch between the designations in a HyperSwap or mirroring relation and the Multi-site roles. Overwrite the existing designations?

## Return codes

- **TARGET_BAD_NAME**

  The target name does not exist.
- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **CONS_GROUP_BAD_NAME**

  The consistency group name does not exist.
- **MULTISITE_BAD_GLOBAL_ID**

  The Multi-site global ID does not exist.
- **MULTISITE_SMASTER_TARGET_NOT_CONNECTED**

  The Multi-site SMaster target is not connected. The operation cannot be completed.
- **MULTISITE_SLAVE_TARGET_NOT_CONNECTED**

  The Multi-site Slave target is not connected. The operation cannot be completed.
- **MULTISITE_SLAVE_MASTER_RELATION_IS_MISSING**

  The relation exists on the Master but not on the Slave.
- **MULTISITE_SMASTER_INVALID_CONFIGURATION**

  The Multi-site SMaster-Master relation configuration is invalid.
- **MULTISITE_ALREADY_DEFINED**

  A Multi-site is already defined for this volume/consistency group.
- **MULTISITE_MASTER_INVALID_CONFIGURATION**

  The Multi-site Master relation configuration is invalid.
- **MULTISITE_DEFINE_FAILED_TO_ROLLBACK_MANUAL_CLEANUP_REQUIRED**

  An error occurred while defining a Multi-site relation, and the system failed to roll back. Please manually clean up Multi-site relations on all systems.
- **MULTISITE_MASTER_INVALID_RELATION_STATE**

  An invalid relation state in the Master.

  **Troubleshooting:** Make sure the existing relations are either both inactive, or one is active/syncronized and the other is inactive.
- **MULTISITE_INIT_RELATION_BAD_UID**

  The relation exists on the Master, but not on the SMaster or Slave.
- **MULTISITE_SMASTER_MASTER_RELATION_IS_MISSING**

  The relation exists on the Master, but not on the SMaster.

- **MULTISITE_SMASTER_INVALID_RELATION_STATE**

  An invalid relation state in the SMaster.

  **Troubleshooting:** The SMaster-Slave asynchronous mirror must be a Standby relation.
- **MULTISITE_INVALID_MASTER_SLAVE_TARGET_CONNECTIVITY**

  Cannot define a Multi-site relation because Master-Slave target connectivity at the Master system is invalid. Try again in a few seconds.
- **MULTISITE_SLAVE_SMASTER_TARGET_MISMATCH**

  The defined SMaster target does not match the defined Slave-SMaster relation target.
- **MULTISITE_SLAVE_INVALID_CONFIGURATION**

  The Multi-site Slave relation configuration is invalid.
- **MULTISITE_SLAVE_INVALID_RELATION_STATE**

  An invalid relation state in the Slave.

  **Troubleshooting:** The SMaster-Slave asynchronous mirror must be a Standby relation.
- **TARGET_BAD_TYPE**

  The target machine is not an XIV machine.
- **MULTISITE_SMASTER_MASTER_TARGET_MISMATCH**

  The defined Master target does not match the defined SMaster-Master relation target.
- **MULTISITE_SLAVE_MASTER_TARGET_MISMATCH**

  The defined Master target does not match the defined Slave-Master relation target.
- **MULTISITE_INVALID_MASTER_SMASTER_TARGET_CONNECTIVITY**

  Cannot define a Multi-site relation because Master-SMaster target connectivity at the Master system is invalid.
- **MULTISITE_SMASTER_SLAVE_TARGET_MISMATCH**

  The defined Slave target does not match the defined SMaster-Slave relation target.
- **VOLUME_HAS_STANDBY_SNAPSHOTS**

  A Multi-site volume has standby snapshots created by a previous process.
- **MULTISITE_INCOMPATIBLE_TARGET_VERSION**

  Multi-site is not supported between the system versions of the specified peers.
- **MULTISITE_NUM_OF_MULTISITES_ON_SLAVE_LIMIT_REACHED**

  The number of Multi-site relations defined on the Slave system exceeded the limit.
- **RELATION_RETRY_OPERATION**

  There is an operation in progress on this relation , please try again your request in a few seconds.

  **Troubleshooting:** Please try again the command in a few seconds.
- **MULTISITE_SMASTER_INCOMPATIBLE_VERSION**

  The software version on the SMaster system does not support Multi-site.
- **MULTISITE_NUM_OF_MULTISITES_ON_SMASTER_LIMIT_REACHED**

  The number of Multi-site relations defined on the SMaster system exceeded the limit.
- **MULTISITE_SLAVE_INCOMPATIBLE_VERSION**

  The software version on the Slave system does not support Multi-site.
- **MULTISITE_NUM_OF_MULTISITES_ON_MASTER_LIMIT_REACHED**

  The number of Multi-site relations defined on the Master system exceeded the limit.

- **MULTISITE_INVALID_MASTER_LOOPBACK_CONFIGURATION**

  Cannot define a Multi-site relation with a loopback. A loopback target detected on the Master system.
- **MULTISITE_INVALID_SMASTER_LOOPBACK_CONFIGURATION**

  Cannot define a Multi-site relation with a loopback. A loopback target detected on the SMaster system.
- **MULTISITE_INVALID_SLAVE_LOOPBACK_CONFIGURATION**

  Cannot define a Multi-site relation with a loopback. A loopback target detected on the Slave system.
- **MULTISITE_MASTER_SMASTER_CLOCK_SKEW_TOO_BIG**

  Command failed because the time difference between the Master and SMaster systems is too big.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support
- **TARGET_NO_ACCESS**

  No access permissions to the secondary machine.
- **CONS_GROUP_MEMBER_VOL_IS_MISSING_A_RELATION**

  A volume in the consistency group is missing the volume-level Multi-site/mirror/HyperSwap relation.
- **MULTISITE_CANNOT_CONTAIN_SYNC_MIRROR**

  A Multi-site relation cannot include a synchronous mirroring relation.

## Activating an asynchronous mirror within a Multi-site relation

Use the **multisite_activate_async_mirror** command to activate an asynchronous mirror as part of a Multi-site relation.

```
multisite_activate_async_mirror  < vol=VolName | cg=cgName >
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **vol** | Object name | A volume name. | N |
| **cg** | Object name | Master consistency group name. | N |

The command can be issued only on the Master volume or consistency group. Issuing it on a standalone Master, that is, on a peer that originally was a Slave and then was promoted to the Master in a failure recovery procedure, is not allowed.

The **multisite_activate_async_mirror** command first ensures that the asynchronous mirror activation will maintain the consistency of the recovery image on the Slave, and then changes the asynchronous mirror state to Active.

If during the preliminary check the command detects that the last consistent snapshot on the Slave is more recent than on the Master, it emits the MULTISITE_SLAVE_HAS_MORE_RECENT_DATA return code and fails to activate the asynchronous mirror. In this case, the user may proceed as follows:

- Delete the Multi-site relation, and overwrite the more recent data on the Slave with the Master data.
- Delete the Multi-site relation, and restore the Master with the Slave data.
- Switch roles between Master and SMaster, in case the SMaster has more recent data that matches the Slave data.

## Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Allowed |
| Application administrator | Allowed |
| Security administrator | Disallowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

## Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **CONS_GROUP_BAD_NAME**

  The consistency group name does not exist.
- **CONS_GROUP_IS_NOT_MULTISITE**

  The local consistency group does not have Multi-site definitions.
- **VOLUME_IS_NOT_MULTISITE**

  This volume is not part of Multi-site relation.
- **SYNC_ALREADY_ACTIVE**

  Synchronization is already active.
- **MULTISITE_ROLE_IS_NOT_MASTER**

  The local peer is not the Master in the Multi-site relation.
- **MIRROR_CONFIGURATION_ERROR**

  The mirror's local configuration does not match its remote configuration.
- **MULTISITE_CONS_GROUP_MEMBERSHIP_MISMATCH**

  A Multi-site consistency group contains different volumes on Master and Slave. This problem may occur because a previously issued cg_add_vol or cg_remove_vol command completed only partially.

  **Troubleshooting:** Please retry the cg_add_vol or cg_remove_vol to ensure the membership in the Multi-site consistency groups in all 3 peers are equivalent
- **MIRROR_RETRY_OPERATION**

  There is an operation in progress on this mirror.

  **Troubleshooting:** Retry the command in a few seconds.
- **MIRROR_SIZE_MISMATCH**

  The secondary and primary volume sizes are different.
- **REMOTE_MAX_VOLUMES_REACHED**

  The maximum number of volumes on the remote machine is already reached.
- **MAX_VOLUMES_REACHED**

  The maximum allowed number of volumes is already reached.
- **DOMAIN_MAX_VOLUMES_REACHED**

  The domain exceeds the maximum allowed number of volumes.
- **REMOTE_VOLUME_IS_MASTER**

A volume on the remote machine is already defined as primary.

- **MULTISITE_SLAVE_HAS_MORE_RECENT_DATA**

  Cannot activate the Multi-site relation because the Slave has more recent data.
- **REMOTE_MIRROR_IS_STANDBY**

  The remote mirror is marked as Standby.
- **MULTISITE_SLAVE_INVALID_STANDBY_CONFIGURATION**

  The Multi-site Slave standby mirror configuration is invalid.
- **MULTISITE_MASTER_INVALID_CONFIGURATION**

  The Multi-site Master relation configuration is invalid.
- **REMOTE_DOMAIN_MAX_VOLUMES_REACHED**

  The maximum number of volumes in the remote machine domain is already reached.
- **MULTISITE_ASYNC_MIRROR_IS_NOT_CONNECTED**

  Cannot activate a Multi-site relation after role change, if the async mirror is not connected.
- **MULTISITE_MASTER_SLAVE_INCONSISTENT_LRS_CONFIGURATION**

  Cannot activate the Multi-site relation because the last replicated snapshot (LRS) on the Master is more recent than on the Slave. The snapshot configuration is inconsistent.
- **MULTISITE_MASTER_SMASTER_CLOCK_SKEW_TOO_BIG**

  Command failed because the time difference between the Master and SMaster systems is too big.
- **MULTISITE_ROLE_IS_STANDALONE_MASTER**

  Command failed because a Multi-site STANDALONE master cannot be activated.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support
- **REMOTE_DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier of the remote system is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support
- **REMOTE_SYSTEM_OUT_OF_PHYSICAL_SPACE**

  The operation not allowed while the remote system is out of physical space.
- **VOLUME_BELONGS_TO_MULTISITE_CONS_GROUP**

  A Multi-site volume is part of a Multi-site consistency group.
- **VOLUME_BELONGS_TO_MIRRORED_CONS_GROUP**

  The volume mirror is part of a consistency group mirror.
- **MIRROR_CAN_NOT_BE_ACTIVATED**

  Mirroring cannot be activated.
- **CONS_GROUP_MEMBER_VOL_IS_MISSING_A_RELATION**

  A volume in the consistency group is missing the volume-level Multi-site/mirror/HyperSwap relation.
- **REMOTE_MULTISITE_IS_NOT_SMASTER**

  The remote peer is not the SMaster.
- **MULTISITE_SMASTER_INVALID_CONFIGURATION**

  The Multi-site SMaster-Master relation configuration is invalid.
- **MULTISITE_BAD_GLOBAL_ID**

The Multi-site global ID does not exist.

## Deleting a Multi-site relation

Use the **multisite_delete** command to delete an existing Multi-site relation.

```
multisite_delete < vol=VolName | cg=cgName > [ force=<yes|no> ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **vol** | Object name | A volume name. | N | N/A |
| **cg** | Object name | Master consistency group name. | N | N/A |
| **force** | Boolean | Forces the deletion of a Multi-site relation on the local system. | N | no |

This command deletes only the Multi-site relation, the two-way relations that make part of it remain intact.

In a healthy configuration, the **multisite_delete** command is issued on the Master volume or consistency group, and affects all three peers. When issued on a consistency group, it automatically deletes the Multi-site relation from each of the member volumes.

If a Multi-site relation cannot be deleted on every peer due a system failure or connectivity problem, it will be deleted partially. This allows deleting a Multi-site configuration when the connectivity between the systems is down.

If a permanent failure has occurred, and the Multi-site configuration cannot be recovered, **multisite_delete** is used to remove the Multi-site relation from the local site, even though the peers may be disconnected. The **force=yes** parameter can is used in this scenario only.

To delete a Multi-site relation, make sure that only one two-way relation from those defined on the target peer is currently active, either HyperSwap or asynchronous mirror.

When a Multi-site relation is deleted, the Multi-site internal snapshots on the SMaster become external and available for the user to manage. If a new Multi-site relation is created for the same volume or consistency group, those snapshots must be deleted.

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Allowed |
| Application administrator | Disallowed |
| Security administrator | Disallowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

### Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **CONS_GROUP_BAD_NAME**

The consistency group name does not exist.

- **CONS_GROUP_IS_NOT_MULTISITE**

  The local consistency group does not have Multi-site definitions.
- **VOLUME_IS_NOT_MULTISITE**

  This volume is not part of Multi-site relation.
- **MULTISITE_BAD_GLOBAL_ID**

  The Multi-site global ID does not exist.
- **MULTISITE_TOO_MANY_ACTIVE_RELATIONS**

  Only one mirror/HyperSwap relation can be active when a Multi-site relation is being deleted.
- **MULTISITE_ROLE_IS_NOT_MASTER**

  The local peer is not the Master in the Multi-site relation.
- **MULTISITE_DELETE_PARTIAL_FAILURE_MANUAL_CLEANUP_REQUIRED**

  Failed to delete the Multi-site relation on all systems. Manual cleanup is required.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support
- **VOLUME_BELONGS_TO_MULTISITE_CONS_GROUP**

  A Multi-site volume is part of a Multi-site consistency group.
- **MIRROR_RETRY_OPERATION**

  There is an operation in progress on this mirror.

  **Troubleshooting:** Retry the command in a few seconds.

## Changing roles in a Multi-site relation

Use the **multisite_change_role** command to change the role of the local Multi-site relation peer.

```
multisite_change_role  < vol=VolName | cg=cgName > new_role=<Master|None|Slave|SMaster>
```

### Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **vol** | Object name | A volume name. | N |
| **cg** | Object name | Master consistency group name. | N |
| **new_role** | Enumeration | Role name of the local Multi-site relation peer. | Y |

This command is used in the following cases:

- As part of the procedure restoring the original configuration after an automatic or manual failover
- In order to perform a manual failover when the automatic failover did not take place
- When the recovery from a DR (disaster recovery) site is needed.

**multisite_change_role** can be issued only for non-operational relations. It is assumed that the command will be issued on all the relevant peers before the relation becomes operational again, so that either the original roles will be restored, or the SMaster and Master will switch roles.

The following role changes are allowed:

- Master to SMaster. For a successful role change from Master to SMaster, the volume can be in any phase, except Initializing. The Master ceases serving host requests, and is set to act as an SMaster and accept replication from the other peer.
- SMaster to Master. For information about the SMaster-to-Master role change, see below.
- Slave to Master. This role change is available only for a standalone Slave, that is not connected to either Master or SMaster. After such Slave becomes Master, it cannot be incorporated into an existing Multi-site relation. Instead, a new Multi-site relation must be defined.

The SMaster-to-Master role change

Preliminary requirements

- Make sure that the original Master is not available and cannot become available while the other peer is acting as the Master. To verify this, run the **multisite_list** command on the original Master and check the value of the **HA I/O Service** output field. Only if the returned value is Unavailable, proceed with issuing the **multisite_change_role** command on the SMaster.
- Stop the applications that use Multi-site volume(s) or consistency group(s). Note that the applications must be stopped completely and not merely paused. This is to make sure that an application does not use any cached state when accessing the volume or consistency group for the first time after the role change.

**Important:**
A failure to comply with either of the above requirements may result in a data integrity issue.

Data consistency on the SMaster

After an SMaster is successfully changed to a Master, the volume (or consistency group) starts accepting requests from hosts. Upon explicit activation, it starts replicating to the other peer (the original Master).

If HyperSwap is interrupted in the middle of the re-synchronization process, the SMaster may be inconsistent. The last consistent image of the SMaster volume (or consistency group) is preserved in the last consistent snapshot (LCS), which is automatically created before the re-synchronization starts. If the LCS exists, the command emits a ARE_YOU_SURE_YOU_WANT_TO_CHANGE_MULTISITE_ROLE_WITH_LCS warning. In this case, before issuing the before issuing the **multisite_change_role** command, the administrator must choose whether to use the existing contents of the original SMaster, which may be inconsistent, or revert the original SMaster to the LCS.

## Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Allowed |
| Application administrator | Disallowed |
| Security administrator | Disallowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

## Warnings

- **SOME_DATA_WILL_BE_LOST_ARE_YOU_SURE**

  Are you sure you want the mirror/HyperSwap local peer to become secondary and lose the data that was not replicated?

- **ARE_YOU_SURE_YOU_WANT_TO_CHANGE_MULTISITE_ROLE**

  Are you sure you want to change the Multi-site role?

- **ARE_YOU_SURE_YOU_WANT_TO_CHANGE_MULTISITE_ROLE_WITH_LCS**

  Are you sure you want to change the Multi-site role with the existing last consistent snapshot?

## Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.
- **CONS_GROUP_BAD_NAME**

  The consistency group name does not exist.
- **CONS_GROUP_IS_NOT_MULTISITE**

  The local consistency group does not have Multi-site definitions.
- **VOLUME_IS_NOT_MULTISITE**

  This volume is not part of Multi-site relation.
- **MULTISITE_SMASTER_CAN_ONLY_BE_CHANGED_TO_MASTER**

  On this system, the Multi-site SMaster can be changed only to the Master.
- **MULTISITE_SLAVE_CAN_ONLY_BE_CHANGED_TO_MASTER**

  On this system, the Multi-site Slave can be changed only to the Master.
- **MULTISITE_MASTER_CAN_ONLY_BE_CHANGED_TO_SLAVE**

  On this system, the Multi-site Master can be changed only to the Slave.
- **MIRROR_RETRY_OPERATION**

  There is an operation in progress on this mirror.

  **Troubleshooting:** Retry the command in a few seconds.
- **HA_RETRY_OPERATION**

  An operation is in progress on this HyperSwap relation.

  **Troubleshooting:** Try issuing the command again in a few seconds.
- **MULTISITE_MASTER_CAN_ONLY_BE_CHANGED_TO_SMASTER**

  On this system, the Multi-site Master can be changed only to the SMaster.
- **MIRROR_IS_INITIAL**

  The operation is not permitted during the Initialization phase.
- **HA_IS_INITIAL**

  The operation is not permitted during the HyperSwap relation initialization phase.
- **MIRROR_HAS_NO_SYNCHED_SNAPSHOT**

  The mirror does not have a synchronized snapshot.
- **HA_HAS_NO_SYNCHED_SNAPSHOT**

  This HyperSwap volume does not have a synchronized snapshot.
- **MASTER_CANNOT_BE_DEMOTED**

  The primary volume cannot be demoted to secondary. Peer status mismatch.
- **MIRROR_IS_ACTIVE**

  Remote mirroring is currently active.
- **HA_IS_ACTIVE**

  HyperSwap is currently active.
- **VOLUME_HAS_DATA_MIGRATION**

Data Migration is defined for this volume.

- **COMMAND_NOT_SUPPORTED_FOR_OLVM_VOLUMES**

  This command is not supported for IBM Hyper-Scale Mobility volumes.

- **VOLUME_BELONGS_TO_MULTISITE_CONS_GROUP**

  A Multi-site volume is part of a Multi-site consistency group.

- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support

# Adding a Standby mirror to a Multi-site relation

Use the **multisite_register_standby_mirror** command to register a Standby mirror in an existing Multi-site relation.

```
multisite_register_standby_mirror < vol=VolName | cg=cgName >
```

## Parameters

| Name | Type | Description | Mandatory |
|---|---|---|---|
| **vol** | Object name | A volume name. | N |
| **cg** | Object name | Master consistency group name. | N |

Issuing this command is only allowed on the system that holds the SMaster.

As a preliminary requirement for **multisite_register_standby_mirror**, the user must create the SMaster-Slave asynchronous mirroring relation.

**multisite_register_standby_mirror** verifies the configuration of the existing SMaster-Slave mirroring relation and adds it as a registered Standby to the Multi-site relation.

When this command is applied to a consistency group, the Standby relation is registered on every volume in the consistency group.

## Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Allowed |
| Application administrator | Disallowed |
| Security administrator | Disallowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

## Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.

- **CONS_GROUP_BAD_NAME**

  The consistency group name does not exist.

- **CONS_GROUP_IS_NOT_MULTISITE**

The local consistency group does not have Multi-site definitions.

- **VOLUME_IS_NOT_MULTISITE**

  This volume is not part of Multi-site relation.
- **TARGET_BAD_NAME**

  The target name does not exist.
- **MULTISITE_ROLE_IS_NOT_SMASTER**

  The local peer is not the SMaster in the Multi-site relation.
- **MULTISITE_SMASTER_INVALID_CONFIGURATION**

  The Multi-site SMaster-Master relation configuration is invalid.
- **MULTISITE_SLAVE_INVALID_STANDBY_CONFIGURATION**

  The Multi-site Slave standby mirror configuration is invalid.
- **MULTISITE_STANDBY_RELATION_ALREADY_DEFINED**

  A standby relation is already defined on this system.
- **MULTISITE_SLAVE_SMASTER_TARGET_MISMATCH**

  The defined SMaster target does not match the defined Slave-SMaster relation target.
- **MULTISITE_SLAVE_INVALID_CONFIGURATION**

  The Multi-site Slave relation configuration is invalid.
- **MULTISITE_SLAVE_INVALID_RELATION_STATE**

  An invalid relation state in the Slave.

  **Troubleshooting:** The SMaster-Slave asynchronous mirror must be a Standby relation.
- **MULTISITE_SMASTER_INVALID_RELATION_STATE**

  An invalid relation state in the SMaster.

  **Troubleshooting:** The SMaster-Slave asynchronous mirror must be a Standby relation.
- **MULTISITE_SMASTER_INVALID_STANDBY_CONFIGURATION**

  The standby relation configuration on SMaster is invalid.
- **MULTISITE_SMASTER_SLAVE_TARGET_MISMATCH**

  The defined Slave target does not match the defined SMaster-Slave relation target.
- **MULTISITE_INVALID_SMASTER_LOOPBACK_CONFIGURATION**

  Cannot define a Multi-site relation with a loopback. A loopback target detected on the SMaster system.
- **MULTISITE_INVALID_SLAVE_LOOPBACK_CONFIGURATION**

  Cannot define a Multi-site relation with a loopback. A loopback target detected on the Slave system.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support
- **CONS_GROUP_MEMBER_VOL_IS_MISSING_A_RELATION**

  A volume in the consistency group is missing the volume-level Multi-site/mirror/HyperSwap relation.
- **VOLUME_BELONGS_TO_MULTISITE_CONS_GROUP**

  A Multi-site volume is part of a Multi-site consistency group.

# Switching roles between Master and SMaster

Use the **multisite_switch_roles** command to switch roles between the Master and SMaster volumes or consistency groups.

```
multisite_switch_roles <vol=VolName | cg=cgName> [ verify_standby=<yes|no> ]
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **vol** | Object name | Local volume name. | N | N/A |
| **cg** | Object name | Local consistency group name. | N | N/A |
| **verify_standby** | Boolean | If set to yes, the command can be successfully completed only if the SMaster-Slave asynchronous Standby mirror can be activated and the SMaster can undertake the disaster recovery. If set to no, the command's successful completion is not contingent on the Standby mirror state. | N | no |

This command can only be issued on the Master volume or consistency group. The command can be successfully completed only if:

- the HyperSwap relation is synchronized
- the asynchronous mirroring relation between the Master and the Slave is deactivated or not configured

In addition, if **verify_standby** is set to yes, the command verifies whether the SMaster-Slave asynchronous Standby mirror can be activated and the SMaster can undertake the disaster recovery.

**Important:**

Conditions may change after the verification. As a result, the subsequent activation attempt may still fail even though all conditions were verified before switching the roles.

If all the requirements are met, **multisite_switch_roles** allows the system to perform all pending writes. Only after they have been committed, the roles are switched.

Following the execution of the command:

- The volume or consistency group that was previously the Master becomes the SMaster
- The volume or consistency group that was previously the SMaster becomes the Master
- The Multi-site relation remains active
- The asynchronous mirroring relation between the SMaster and Slave is considered registered standby, provided that it had been active and then deactivated in preparation for the role switch

## Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Allowed |
| Application administrator | Disallowed |

| User Category | Permission |
|---|---|
| Security administrator | Disallowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

## Warnings

- **ARE_YOU_SURE_YOU_WANT_TO_SWITCH_ROLES**

  Are you sure you want to switch the roles in this relation?

- **ARE_YOU_SURE_YOU_WANT_TO_SWITCH_THE_PEER_ROLES**

  The system that owns the primary volume is currently not connected to the Quorum Witness. If you switch the peer roles, the automatic failover may become unavailable. Are you sure you want to continue?

## Return codes

- **VOLUME_BAD_NAME**

  The volume name does not exist.

- **CONS_GROUP_BAD_NAME**

  The consistency group name does not exist.

- **MULTISITE_ROLE_IS_NOT_MASTER**

  The local peer is not the Master in the Multi-site relation.

- **VOLUME_BELONGS_TO_MULTISITE_CONS_GROUP**

  A Multi-site volume is part of a Multi-site consistency group.

- **MULTISITE_STANDBY_MIRROR_NOT_REGISTERED**

  A standby mirror is either not registered or not defined for this Multi-site relation.

- **MULTISITE_STANDBY_MIRROR_IS_NOT_CONNECTED**

  There is currently no connection between the SMaster and the standby async mirror target system.

  **Troubleshooting:** The relation is part of a Multi-site relation, therefore connectivity between the SMaster and the standby async mirror Slave is required

- **MULTISITE_ASYNC_MIRROR_IS_ACTIVE**

  The mirror between the Multi-site Master and Slave is active.

  **Troubleshooting:** The relation in which the roles must be switched is part of a Multi-site relation. The async mirror on the Master must be deactivated.

- **MULTISITE_SMASTER_HAS_NO_SYNCHED_SNAPSHOT**

  The Multi-site SMaster does not have a synchronized snapshot.

- **HA_IS_NOT_SYNCHRONIZED**

  The HyperSwap relation is not synchronized.

- **VOLUME_HAS_DATA_MIGRATION**

  Data Migration is defined for this volume.

- **REMOTE_TARGET_NOT_CONNECTED**

  There is currently no connection from the target system.

- **HA_CONNECTIVITY_NOT_SUFFICIENT**

  The connectivity between the systems is not sufficient for the automatic failover.

- **HA_RELATION_MASTER_COULD_NOT_UPDATE_QW_AFTER_RETURN_TO_GOOD_STATE**

  The primary volume had control over the relationship during a past failure, and was unable to update the Quorum Witness after the recovery.
- **HA_HAS_SYNC_JOB**

  This operation is not permitted on a HyperSwap relation with active sync jobs.
- **HA_RETRY_OPERATION**

  An operation is in progress on this HyperSwap relation.

  **Troubleshooting:** Try issuing the command again in a few seconds.
- **REMOTE_HA_IS_NOT_ACTIVE**

  The remote peer in this HyperSwap relation is not active.
- **COMMAND_NOT_SUPPORTED_FOR_OLVM_VOLUMES**

  This command is not supported for IBM Hyper-Scale Mobility volumes.
- **DATA_REDUCTION_TIER_IS_OFFLINE**

  The data reduced tier is offline, the operation is not allowed.

  **Troubleshooting:** Contact IBM Support
- **SYSTEM_OUT_OF_PHYSICAL_SPACE**

  The operation not allowed while the system is out of physical space.
- **VOLUME_IS_NOT_MULTISITE**

  This volume is not part of Multi-site relation.
- **CONS_GROUP_IS_NOT_MULTISITE**

  The local consistency group does not have Multi-site definitions.
- **VOLUME_BELONGS_TO_HA_CONS_GROUP**

  This volume belongs to a HyperSwap consistency group. The operation can be carried out on a volume which is not part of a consistency group, or on the consistency group itself.
- **REMOTE_MAY_NOT_HAVE_COMPLETED_THE_OPERATION**

  The operation may be not yet completed on the remote target.
- **MIRROR_RETRY_OPERATION**

  There is an operation in progress on this mirror.

  **Troubleshooting:** Retry the command in a few seconds.

# Chapter 26. Security configuration commands

This chapter describes the command line interface (CLI) for security configuration.

## Listing configuration parameters for a communication protocol

Use the **protocol_config_list** command to list configuration parameters per communication protocol.

```
protocol_config_list [ protocol=<xcli|kmip|cim> ]
```

### Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **protocol** | Enumeration | The available options are: XCLI, KMIP, CIM, or All (if no value is specified). | N | All (if no value is specified). |

The following default parameters are shown:

- Protocol Name (XCLI, KMIP, CIM)
- Minimal TLS Version (TLS1.0, TLS1.1, TLS1.2)

**Example:**

```
protocol_config_list
```

**Output:**

```
Protocol Name            Minimal TLS Version
------------------------ ----------------------------------
CIM                      TLS1.2
KMIP                     TLS1.2
XCLI                     TLS1.2
```

| Field ID | Field output | Default position |
|---|---|---|
| **protocol_name** | Protocol Name | 1 |
| **min_tls_level** | Minimal TLS Version | 2 |

### Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Allowed |
| Security administrator | Allowed |
| Read-only users | Allowed |
| Technicians | Allowed |

© Copyright IBM Corp. 2016, 2022

# Setting configuration parameters for a communication protocol

Use the **protocol_config_set** command to set configuration parameters for a specific communication protocol or for all protocols.

```
protocol_config_set [ protocol=<xcli|kmip|cim> ] min_tls_level=<tls1.0|tls1.1|tls1.2>
```

## Parameters

| Name | Type | Description | Mandatory | Default |
|---|---|---|---|---|
| **protocol** | Enumeration | The available options are: XCLI, KMIP, CIM (case insensitive), or All (if no value is specified). | N | All (if no value is specified). |
| **min_tls_level** | Enumeration | Minimum TLS version to support. The available options are: TLS1.0, TLS1.1, or TLS1.2 (case insensitive). | Y | N/A |

**Example:**

```
protocol_config_set protocol=KMIP min_tls_level=TLS1.2
```

**Output:**

```
Command completed successfully
```

## Access control

| User Category | Permission |
|---|---|
| Storage administrator | Allowed |
| Storage integration administrator | Disallowed |
| Application administrator | Disallowed |
| Security administrator | Allowed |
| Read-only users | Disallowed |
| Technicians | Disallowed |

# Chapter 27. Events

## VOLUME_CREATE

| Severity | Description |
|---|---|
| informational | Volume was created with name '*volume.name*' and size *volume.size*GB in Storage Pool with name '*volume.pool_name*'. |

## VOLUME_CREATE_FAILED_TOO_MANY

| Severity | Description | Troubleshooting |
|---|---|---|
| warning | Volume with name '*name*' could not be created. You are attempting to add more volumes than the system permits. | Delete volumes to allow new ones to be created. |

## VOLUME_RENAME

| Severity | Description |
|---|---|
| informational | Volume with name '*old_name*' was renamed '*volume.name*'. |

## SUBORDINATE_VOLUME_RENAME

| Severity | Description |
|---|---|
| informational | Remote volume with name '*old_name*' was renamed '*volume.name*' by local system. |

## VOLUME_RESIZE

| Severity | Description |
|---|---|
| informational | Volume with name '*volume.name*' was resized from *old_size*GB to *volume.size*GB. |

## SECONDARY_VOLUME_RESIZE

| Severity | Description |
|---|---|
| informational | Secondary volume with name '*volume.name*' was resized by primary machine from *old_size*GB to *volume.size*GB. |

## VOLUME_DELETE

| Severity | Description |
|---|---|
| informational | Volume with name '*volume.name*' was deleted and its data is no longer accessible. |

## VOLUME_FORMAT

| Severity | Description |
|---|---|
| informational | Volume with name '*volume.name*' was formatted. |

© Copyright IBM Corp. 2016, 2022

## 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. |

## HA_VOLUME_MAPPING_MISSING

| Severity | Description | Troubleshooting |
|---|---|---|
| major | The HyperSwap volume '*local volume name*' must be mapped to allow a host to continue to perform I/O after an automatic failover. | Map the volume to the relevant host. |

## HA_VOLUME_MAPPING_OK

| Severity | Description |
|---|---|
| informational | The HyperSwap volume '*local volume name*' is now mapped to a host. |

## DATA_REBUILD_STARTED

| Severity | Description |
|---|---|
| informational | Rebuild process started because system data is not protected. *data_percent*% of the data must be rebuilt. |

## DATA_REBUILD_COMPLETED

| Severity | Description |
|---|---|
| informational | Rebuild process completed. System data is now protected. |

## DATA_REDIST_STARTED

| Severity | Description |
|---|---|
| informational | Starting data transfer to new disks. |

## DATA_REDIST_COMPLETED

| Severity | Description |
|---|---|
| informational | Completed data transfer to new disks. |

## DATA_REBUILD_COMPLETED_REDIST_STARTED

| Severity | Description |
|---|---|
| informational | Rebuild process completed. System data is now protected. Starting data transfer to new disks. |

## DATA_REDIST_TIME_LIMIT_EXCEEDED

| Severity | Description |
|---|---|
| warning | Data redistribution is taking too long. *data_percent*% of the required redistribution still remains to be done. |

## DATA_REDIST_BLOCKED

| Severity | Description |
|---|---|
| informational | Blocking data transfer to new phased in media for *delay_in_seconds* seconds. |

## DATA_REDIST_ALLOWED

| Severity | Description |
|---|---|
| informational | Allowing data transfer to new phased in media. |

## STORAGE_POOL_EXHAUSTED

| Severity | Description | Troubleshooting |
|---|---|---|
| major | Pool '*pool*' is full. All volumes are locked. | Enlarge Storage Pool or move or delete volumes or Clones with Clone Deletion Priority 0. |

## STORAGE_POOL_UNLOCKED

| Severity | Description |
|---|---|
| major | Pool '*pool*' has empty space. All volumes are unlocked. |

## STORAGE_POOL_SNAPSHOT_USAGE_INCREASED

| Severity | Description |
|---|---|
| variable | Usage by snapshots of Storage Pool with name '*pool.name*' has reached *current*%. |

## STORAGE_POOL_SNAPSHOT_USAGE_DECREASED

| Severity | Description |
|---|---|
| informational | Usage by snapshots of Storage Pool with name '*pool.name*' has decreased to *current*%. |

## HOST_CONNECTED

| Severity | Description |
|---|---|
| informational | Host '*host*' has connected to the system. |

## HOST_DISCONNECTED

| Severity | Description |
|---|---|
| warning | Host '*host*' has disconnected from the system. |

## HOST_MULTIPATH_OK

| Severity | Description |
|---|---|
| informational | Host '*host*' has redundant connections to the system. #paths=*npaths* |

## HOST_NO_MULTIPATH_ONLY_ONE_PORT

| Severity | Description |
|---|---|
| warning | Host '*host*' is connected to the system through only one of its ports. #paths=*npaths* |

## HOST_MULTIPATH_NOT_RESILIENT_TO_MODULES_FAILURES

| Severity | Description | Troubleshooting |
|---|---|---|
| informational | Host '*host*' is connected to the system through less than three modules and hence is less resilient to modules failures. | Add host paths to the system through at least three active modules. |

## HOST_NO_MULTIPATH_ONLY_ONE_MODULE

| Severity | Description |
|---|---|
| informational | Host '*host*' is connected to the system through only one interface module. #paths=*npaths* |

## SYSTEM_CAPACITY_USAGE_INCREASED

| Severity | Description |
|---|---|
| variable | Usage of system physical capacity increased to *current*%. |

## SYSTEM_CAPACITY_USAGE_DECREASED

| Severity | Description |
|---|---|
| informational | Usage of system physical capacity decreased to *current*%. |

## POOL_CREATE

| Severity | Description |
|---|---|
| informational | Storage Pool of size *pool.size*GB*sparse_type* was created with name '*pool.name*'. |

## POOL_CREATE_FAILED_TOO_MANY

| Severity | Description | Troubleshooting |
|---|---|---|
| warning | Storage Pool with name '*name*' could not be created. You are attempting to add more Storage Pools than the system permits. | Delete Storage Pools to allow new ones to be created. |

## POOL_RENAME

| Severity | Description |
|---|---|
| informational | Storage Pool with name '*old_name*' was renamed '*pool.name*'. |

## POOL_RESIZE

| Severity | Description |
|---|---|
| informational | Storage Pool with name '*pool.name*' was resized from size *old_size*GB*old_sparse_type* to *pool.size*GB*sparse_type*. |

## POOL_RESIZE_SNAPSHOTS

| Severity | Description |
|---|---|
| informational | Snapshot size of Storage Pool with name '*pool.name*' was resized from size *old_size*GB to *pool.snapshot_size*GB. |

## POOL_CHANGE_LOCK_BEHAVIOR

| Severity | Description |
|---|---|
| informational | Lock Behavior of Storage Pool with name '*pool.name*' is now '*state*'. |

## POOL_CONVERTED_TO_SPARSE

| Severity | Description |
|---|---|
| informational | Storage Pool with name '*pool.name*' was converted to sparse. |

## POOL_CONVERTED_TO_REGULAR

| Severity | Description |
|---|---|
| informational | Storage Pool with name '*pool.name*' was converted to regular. |

## POOL_CHANGE_PERF_CLASS

| Severity | Description |
|---|---|
| informational | Performance Class of Storage Pool with name '*pool.name*' is now '*pool.perf_class*'. |

## POOL_CONFIG_SNAPSHOTS

| Severity | Description |
|---|---|
| informational | Management policy of Mirroring snapshots of Storage Pool with name '*pool.name*' has changed'. |

## POOL_DELETE

| Severity | Description |
|---|---|
| informational | Storage Pool with name '*pool.name*' was deleted. |

## FLASH_VDISK_TOO_SMALL

| Severity | Description |
|---|---|
| major | Flash vdisk *name* on *Flash Enclosure* is too small. |

## FLASH_VDISK_LARGER_THAN_EXPECTED

| Severity | Description |
|---|---|
| warning | Flash vdisk *name* on *Flash Enclosure* is larger than expected. |

## DATA_REDUCTION_TIER_IS_OFFLINE

| Severity | Description |
|---|---|
| critical | Data reduction tier moved to offline mode |

## DATA_REDUCTION_RECOVERY_STARTED

| Severity | Description |
|---|---|
| Informational | Data reduction recovery process started |

## DATA_REDUCTION_RECOVERY_FINISHED

| Severity | Description |
|---|---|
| Informational | Data reduction recovery process ended |

## DATA_REDUCTION_RECOVERY_FAILED

| Severity | Description |
|---|---|
| Minor | Data reduction recovery process failed with reason: '*reason*' |

## DATA_REDUCTION_RECOVERY_ABORT_STARTED

| Severity | Description |
|---|---|
| Informational | Data reduction recovery abort started |

## DATA_REDUCTION_RECOVERY_ABORT_FINISHED

| Severity | Description |
|---|---|
| Informational | Data reduction recovery abort finished |

## DATA_REDUCTION_TIER_IS_ONLINE

| Severity | Description |
|---|---|
| informational | Data reduction tier moved to online state |

## DATA_REDUCTION_TIER_BECOMING_OFFLINE

| Severity | Description |
|---|---|
| major | Data reduction tier is transitioning to offline mode |

## DATA_REDUCTION_RESUME_ONLINE_START

| Severity | Description |
|---|---|
| informational | Data reduction is resuming from offline mode |

## DATA_REDUCTION_RESUME_ONLINE_FAILED

| Severity | Description |
|---|---|
| warning | Data reduction failed resuming from offline mode |

## DATA_REDUCTION_COMPRESSION_ADAPTER_FAILED

| Severity | Description |
|---|---|
| Major | Data Reduction node *reporting_node* reported about compression adapter failure and will be killed. |

## FLASH_COMPONENT_INITIALIZING

| Severity | Description | Troubleshooting |
|---|---|---|
| informational | *Flash Component ID* initializing. | Contact IBM Support |

## FLASH_COMPONENT_OK

| Severity | Description | Troubleshooting |
|---|---|---|
| informational | *Flash Component ID* status ok. | Contact IBM Support |

## FLASH_COMPONENT_FAILED

| Severity | Description | Troubleshooting |
|---|---|---|
| variable | *Flash Component ID* has failed. | Contact IBM Support |

## FLASH_COMPONENT_REPLACED

| Severity | Description | Troubleshooting |
|---|---|---|
| informational | *Component ID* was replaced. New serial *New Serial* Old serial *Old Serial* | Contact IBM Support |

## FLASH_FW_HOT_UPGRADE_STARTED

| Severity | Description |
|---|---|
| informational | Start upgrade to version *version*. Enclosure id *component id* |

## FLASH_FW_HOT_UPGRADE_FINISHED

| Severity | Description |
|---|---|
| informational | Finished upgrade to version *version*. Enclosure id *component id* |

## FLASH_UPGRADE_RESUMED

| Severity | Description |
|---|---|
| informational | Finished upgrade resumed. Enclosure id *component id* |

## FLASH_FW_HOT_UPGRADE_RESUMED

| Severity | Description |
|---|---|
| informational | Finished upgrade resumed. Enclosure id *component id* |

## FLASH_UPGRADE_STOPPED

| Severity | Description |
|---|---|
| major | Flash upgrade stopped after *percents*%, reason: *reason*. Enclosure id *component id* |

## FLASH_FW_HOT_UPGRADE_STOPPED

| Severity | Description |
|---|---|
| major | Flash upgrade stopped after *percents*%, reason: *reason*. Enclosure id *component id* |

## FLASH_FW_HOT_UPGRADE_FAILED

| Severity | Description |
|---|---|
| major | Flash firmware hot upgrade failed, reason: *reason*. error: *ccl_error*. Enclosure id *component id* |

## FLASH_FW_HOT_UPGRADE_TIMEOUT

| Severity | Description |
|---|---|
| major | Timeout while upgrading *component id*, progress: *percents* |

## FLASH_RAID_STATUS_CHANGED

| Severity | Description |
|---|---|
| variable | Flash array raid status changed to *status*. Enclosure id *component id*. |

## FLASH_ARRAY_STATUS_CHANGED

| Severity | Description |
|---|---|
| variable | Flash array status changed to *status*. Enclosure id *component id*. |

## FLASH_ENCRYPTION_STATUS_CHANGED

| Severity | Description |
|---|---|
| informational | Encryption enabled changed to *encrypted*. Enclosure id *component id*. |

## FLASH_ENCRYPTION_ENABLE_NOT_ALLOWED

| Severity | Description |
|---|---|
| major | Encryption enable not allowed: *reason*. Enclosure id *component id*. |

## FLASH_ENCRYPTION_DISABLE_NOT_ALLOWED

| Severity | Description |
|---|---|
| major | Encryption disable not allowed: *reason*. Enclosure id *component id*. |

## ENCRYPT_ENABLE_FLASH_ENCLOSURE_FAILED

| Severity | Description |
|---|---|
| major | Encryption enable failed for *component id*: *reason* |

## ENCRYPT_DISABLE_FLASH_ENCLOSURE_FAILED

| Severity | Description |
|---|---|
| major | Encryption disable failed for *component id*: *reason* |

## FLASH_ENCRYPTION_UNLOCK_FAILED

| Severity | Description |
|---|---|
| major | Flash Encryption unlock failed. Enclosure: *component id*. |

## FLASH_ENCLOSURE_WIPEOUT_FAILED

| Severity | Description |
|---|---|
| major | Wipeout failed for *component id*: *reason* |

## FLASH_BBU_CHARGING_STATUS_CHANGED

| Severity | Description |
|---|---|
| informational | BBU charging status changed to *status*. BBU: *component id*. |

## FLASH_BBU_CALIBRATION_STARTED

| Severity | Description |
|---|---|
| informational | BBU calibration started, BBU: *component id*. |

## FLASH_BBU_CALIBRATION_STOPPED

| Severity | Description |
|---|---|
| informational | BBU calibration stopped, BBU: *component id*. |

## FLASH_BBU_CALIBRATION_FAILED

| Severity | Description |
|---|---|
| minor | BBU calibration failed, BBU: *component id*. |

## FLASH_CANISTER_CONNECTED_VIA_SERIAL_CABLE

| Severity | Description |
|---|---|
| informational | Established serial connection with *component id*. |

## FLASH_CANISTER_CONNECTION_VIA_SERIAL_OK

| Severity | Description |
|---|---|
| informational | Established serial connection with *component id*. |

## NO_CONNECTION_TO_FLASH_CANISTER_VIA_SERIAL

| Severity | Description | Troubleshooting |
|---|---|---|
| warning | Failed to connect to *component id* via serial. | Contact IBM Support |

## FLASH_CANISTER_NO_CONNECTION_VIA_SERIAL

| Severity | Description | Troubleshooting |
|---|---|---|
| warning | Failed to connect to *component id* via serial. | Contact IBM Support |

## NO_CONNECTION_TO_FLASH_CANISTER_VIA_ETH

| Severity | Description | Troubleshooting |
|---|---|---|
| warning | Failed to connect to *component id* via ethernet. | Contact IBM Support |

## FLASH_CANISTER_NO_CONNECTION_VIA_ETH

| Severity | Description | Troubleshooting |
|---|---|---|
| warning | Failed to connect to *component id* via ethernet. | Contact IBM Support |

## FLASH_CANISTER_CONNECTION_VIA_ETH_OK

| Severity | Description | Troubleshooting |
|---|---|---|
| informational | Established connection to *component id* via ethernet. | Contact IBM Support |

## FLASH_CANISTER_IP_CHANGED

| Severity | Description | Troubleshooting |
|---|---|---|
| informational | Changed connected ip of *component id* to *IP*. | Contact IBM Support |

## NO_CONNECTION_TO_FLASH_CANISTER_VIA_GW

| Severity | Description | Troubleshooting |
|---|---|---|
| warning | Failed to connect to *component id* via *module id*. | Contact IBM Support |

## FLASH_CANISTER_NO_CONNECTION_VIA_GW

| Severity | Description | Troubleshooting |
|---|---|---|
| warning | Failed to connect to *component id* via *module id*. | Contact IBM Support |

## FLASH_CANISTER_CONNECTED_VIA_GATEWAY

| Severity | Description | Troubleshooting |
|---|---|---|
| informational | Established connection to *component id* via *module id*. | Contact IBM Support |

## FLASH_CANISTER_CONNECTION_VIA_GW_OK

| Severity | Description | Troubleshooting |
|---|---|---|
| informational | Established connection to *component id* via *module id*. | Contact IBM Support |

## FLASH_CANISTER_ETH_LINK_MISWIRE

| Severity | Description | Troubleshooting |
|---|---|---|
| minor | Flash canister *component id* is miswired. | Contact IBM Support |

## FLASH_CANISTER_ETH_LINK_WIRING_OK

| Severity | Description | Troubleshooting |
|---|---|---|
| informational | Flash canister *component id* is no longer miswired. | Contact IBM Support |

## FLASH_CANISTER_GET_CONF_FAILED

| Severity | Description | Troubleshooting |
|---|---|---|
| minor | Failed to get the configuration from *component id*. | Contact IBM Support |

## FLASH_CANISTER_GET_CONF_OK

| Severity | Description | Troubleshooting |
|---|---|---|
| informational | Succeeded to get the configuration from *component id*. | Contact IBM Support |

## FLASH_CANISTER_IN_SERVICE_MODE

| Severity | Description | Troubleshooting |
|---|---|---|
| major | Flash canister is in service mode *service*, canister: *component id*. | Contact IBM Support |

## FLASH_CANISTER_NO_LONGER_IN_SERVICE_MODE

| Severity | Description |
|---|---|
| informational | Flash canister is no longer in service mode: *component id*. |

## FLASH_COMPONENT_TEMPERATURE_OK

| Severity | Description |
|---|---|
| informational | Flash component *component id* temperature is within allowed limits. |

## FLASH_COMPONENT_TEMPERATURE_ABOVE_NORMAL

| Severity | Description |
|---|---|
| warning | Flash component *component id* temperature is above normal. |

## FLASH_COMPONENT_TEMPERATURE_HIGH

| Severity | Description |
|---|---|
| major | Flash component *component id* temperature is high, it exceeds operational level. |

## FLASH_COMPONENT_TEMPERATURE_CRITICALLY_HIGH

| Severity | Description |
|---|---|
| critical | Flash component *component id* temperature is critical. |

## FLASH_ENCLOSURE_STATUS_CHANGED

| Severity | Description | Troubleshooting |
|---|---|---|
| informational | *component id* status changed to *New Status*. | Contact IBM Support |

## FLASH_PSU_FAN_FAILED

| Severity | Description |
|---|---|
| minor | *component id* fan has failed. |

## FLASH_PSU_FAN_OK

| Severity | Description |
|---|---|
| informational | *component id* is ok. |

## FLASH_PSU_HAS_NO_INPUT_POWER

| Severity | Description |
|---|---|
| major | PSU power supply has no input (AC) power. PSU id *component id*. |

## FLASH_PSU_HAS_INPUT_POWER

| Severity | Description |
|---|---|
| informational | PSU power supply has input (AC) power. PSU id *component id*. |

## FLASH_ENCLOSURE_VERSION_IS_UNEXPECTED

| Severity | Description |
|---|---|
| major | *component id* version is *version*, expected version is *expected*. |

## FLASH_ENCLOSURE_NEWER_VERSION_EXISTS

| Severity | Description |
|---|---|
| minor | *component id* version *version* is supported. newer version *expected* exists. |

## FLASH_PSU_COMMUNICATION_ERROR

| Severity | Description |
|---|---|
| minor | PSU communication error. PSU id *component id*. |

## FLASH_BBU_VPD_IS_NOT_VALID

| Severity | Description |
|---|---|
| major | BBU VPD is not valid. BBU id *component id*. |

## FLASH_CARD_UNSUPPORTED

| Severity | Description |
|---|---|
| warning | Flash card is unsupported. Flash card id *component id*. |

## FLASH_BBU_END_OF_LIFE

| Severity | Description |
|---|---|
| warning | Battery is at end of life. BBU id *component id*. |

## FLASH_BBU_NEARING_END_OF_LIFE

| Severity | Description |
|---|---|
| warning | Battery is nearing end of life. BBU id *component id*. |

## FLASH_CARD_COMMUNICATION_ERROR

| Severity | Description |
|---|---|
| minor | Flash card communication error. Flash card id *component id*. |

## FLASH_FAN_COMMUNICATION_ERROR

| Severity | Description |
|---|---|
| minor | Fan communication error. Fan: *component id*. |

## FLASH_ENCLOSURE_THERMAL_THREASHOLD_EXCEEDED

| Severity | Description |
|---|---|
| critical | Flash enclosure *component id* has shutdown after exceeding the thermal threshold. |

## FLASH_ENCLOSURE_ARRAY_OFFLINE

| Severity | Description |
|---|---|
| critical | Flash enclosure *component id* array is offline. |

## FLASH_ENCLOSURE_STARTED_PHASEOUT

| Severity | Description |
|---|---|
| informational | System started phasing out *Component ID*. |

## FLASH_ENCLOSURE_FINISHED_PHASEOUT

| Severity | Description |
|---|---|
| informational | System finished phasing out *Component ID*. |

## FLASH_ENCLOSURE_STARTED_PHASEIN

| Severity | Description |
|---|---|
| informational | System started phasing in *Component ID*. |

## FLASH_ENCLOSURE_FINISHED_PHASEIN

| Severity | Description |
|---|---|
| informational | System finished phasing in *Component ID*. |

## FLASH_CR_KEY_SETUP_FAILED

| Severity | Description |
|---|---|
| major | Failed to set challenge-response key on '*Component ID*'. |

## FLASH_CR_KEY_SETUP_OK

| Severity | Description |
|---|---|
| informational | Challenge-response key was successfully set on '*Component ID*'. |

## FLASH_CR_KEY_SETUP_STARTED

| Severity | Description |
|---|---|
| informational | Challenge-response key set started on '*Component ID*'. |

## FC_PORT_HAS_FAILED

| Severity | Description | Troubleshooting |
|---|---|---|
| major | FC port service *port* has failed due to *code* (attempt number *Number of retries*) | Contact IBM Support |

## NTP_SERVER_TIME_DIFFERENCE_TOO_BIG

| Severity | Description | Troubleshooting |
|---|---|---|
| warning | NTP server *NTP Server* sent a transaction with time difference of *Delta* seconds which exceeds the maximal difference of *Max Allowed* seconds. Transaction will be ignored, please check NTP server's and system's times. | Please contact your Administrator. |

## IPSEC_TUNNEL_OPENED

| Severity | Description |
|---|---|
| informational | The IPSec tunnel named '*name*' between module *Module* and *Right IP* was opened |

## IPSEC_TUNNEL_CLOSED

| Severity | Description |
|---|---|
| informational | The IPSec tunnel named '*name*' between module *Module* and *Right IP* was closed |

## IP_ACCESS_CANNOT_RESOLVE_ADDRESS

| Severity | Description |
|---|---|
| informational | Cannot resolve address '*address*' added to the IP access group *IP access group name*. |

## IP_ACCESS_FAILED_SETTING_RULES

| Severity | Description |
|---|---|
| informational | Failed setting IP access rules. |

## USB_ETHERNET_INTERFACE_OK

| Severity | Description |
|---|---|
| informational | USB ethernet interface on module *Module* was reset successfully and is now OK. |

## USB_ETHERNET_INTERFACE_IS_STILL_DOWN

| Severity | Description | Troubleshooting |
|---|---|---|
| minor | USB ethernet interface on module *Module* is down and failed to reset. | Contact IBM Support |

## USB_ETHERNET_INTERFACE_FAILED

| Severity | Description | Troubleshooting |
|---|---|---|
| warning | USB ethernet interface on module *Module* failed. | Contact IBM Support |

## MIRROR_CREATE

| Severity | Description |
|---|---|
| informational | A remote mirror was defined for Volume '*local volume name*'on Target '*target name*'. Remote Volume is '*remote volume name*'. |

## CG_MIRROR_CREATE

| Severity | Description |
|---|---|
| informational | A remote mirror was defined for Consistency Group '*local CG name*' on Target '*target name*'. Remote Consistency Group is '*remote CG name*'. |

## MIRROR_CREATE_SLAVE

| Severity | Description |
|---|---|
| informational | A remote mirror was defined by Target '*target name*' for Volume '*local volume name*'. Remote Volume is '*remote volume name*'. |

## CG_MIRROR_CREATE_SLAVE

| Severity | Description |
|---|---|
| informational | A remote mirror was defined by Target '*target name*' for CG '*local CG name*'. Remote CG is '*remote CG name*'. |

## MIRROR_SCHEDULE_CHANGE

| Severity | Description |
|---|---|
| informational | Schedule of remote mirror of '*local peer name*' is now '*schedule name*'. |

## MIRROR_CREATE_FAILED_TARGET_NOT_CONNECTED

| Severity | Description | Troubleshooting |
|---|---|---|
| warning | Target could not be reached. Target with name '*target.name*' is currently not connected. | Connect the target system to this system. |

## REMOTE_OPERATION_FAILED_TIMED_OUT

| Severity | Description | Troubleshooting |
|---|---|---|
| warning | Operation on remote machine timed out. Invoking '*Function Name*' on target '*Target Name*' timed out. | Retry operation. If problem persists contact IBM support. |

## MIRROR_RESYNC_FAILED

| Severity | Description | Troubleshooting |
|---|---|---|
| major | Synchronization of meta data with mirror failed. Configuration of remote mirror of volume '*local volume name*' on target '*target name*' does not match local configuration. | Make sure configuration on both machines is compatible and activate the mirror. If problem persists contact IBM support. |

## MIRROR_RESYNC_FAILED_DUE_TO_THIN_PROVISIONING

| Severity | Description | Troubleshooting |
|---|---|---|
| major | Synchronization of bitmaps with mirror failed. Not enough hard capacity left in Pool of volume '*mirror.local_volume_name*'. | Delete unnecessary volumes in pool or enlarge the pool's hard size. |

## MIRROR_SYNC_STARTED

| Severity | Description |
|---|---|
| informational | Synchronization of remote mirror of volume '*local volume name*' on Target '*target name*' has started. |

## MIRROR_SYNC_ENDED

| Severity | Description |
|---|---|
| informational | Synchronization of remote mirror of peer '*local peer name*' on target '*target name*' has ended. |

## MIRROR_CANNOT_CREATE_SYNC_JOB_TOO_MANY_VOLUMES

| Severity | Description |
|---|---|
| major | Synchronization of remote mirror of peer '*local peer name*' on target '*target name*' can not be synced , insufficent volume available for this operation. |

## MIRROR_CANNOT_CREATE_LRS_TOO_MANY_VOLUMES

| Severity | Description |
|---|---|
| major | Synchronization of remote mirror of peer '*local peer name*' on target '*target name*' can not be synced , insufficent volume available for this operation. |

## MIRROR_REESTABLISH_FAILED_CONFIGURATION_ERROR

| Severity | Description | Troubleshooting |
|---|---|---|
| major | Mirror reestablish failed. Either configuration of remote mirror of peer '*local peer name*' on target '*target name*' does not match local configuration. | Make sure configuration on both machines is compatible and activate the mirror. If problem persists contact IBM support. |

## MIRROR_ACTIVATE

| Severity | Description |
|---|---|
| informational | The Remote Mirror of peer '*local peer name*' on Target '*target name*' was activated. |

## MIRROR_DEACTIVATE

| Severity | Description |
|---|---|
| informational | The Remote Mirror of peer '*local peer name*' on Target '*target name*' was deactivated. |

## MIRROR_SLAVE_ACTIVATE

| Severity | Description |
|---|---|
| informational | The mirror of peer '*local peer name*' on target '*target name*' was activated. |

## MIRROR_SLAVE_DEACTIVATE

| Severity | Description |
|---|---|
| informational | The mirror of peer '*local peer name*' on target '*target name*' was deactivated. |

## MIRROR_DEACTIVATE_SECONDARY_LOCKED

| Severity | Description |
|---|---|
| minor | The Remote Mirror of peer '*local peer name*' on Target '*target name*' was deactivated since the Pool on the secondary machine was locked. |

## MIRROR_DEACTIVATE_CONFIGURATION_ERROR

| Severity | Description |
|---|---|
| minor | The Remote Mirror of peer '*local peer name*' on Target '*target name*' was deactivated since the Mirror configuration on the slave machine has changed. |

## MIRROR_DELETE

| Severity | Description |
|---|---|
| informational | The Remote Mirror relation of peer '*local peer name*' to a peer on Target '*target name*' was deleted. |

## MIRROR_REVERSE_ROLE_TO_SLAVE

| Severity | Description |
|---|---|
| informational | Local peer '*local peer name*' is now Slave of a peer on Target '*target name*'. |

## MIRROR_REVERSE_ROLE_TO_MASTER

| Severity | Description |
|---|---|
| informational | Local peer '*local peer name*' is now Master of a peer on Target '*target name*'. |

## MIRROR_REVERSE_ROLE_OF_PEER_WITH_LCS_TO_MASTER

| Severity | Description |
|---|---|
| informational | Local peer '*local peer name*' is now Master of a peer on Target '*target name*' external last consistent snapshot should be deleted manually . |

## MIRROR_SWITCH_ROLES_TO_SLAVE

| Severity | Description |
|---|---|
| informational | Local peer '*local peer name*' switched roles with peer on Target '*target name*'. It is now Slave. |

## MIRROR_SWITCH_ROLES_TO_MASTER

| Severity | Description |
|---|---|
| informational | Local peer '*local peer name*' switched roles with peer on Target '*target name*'. It is now Master. |

## MIRROR_REESTABLISH_FAILED_TOO_MANY_VOLUMES

| Severity | Description | Troubleshooting |
|---|---|---|
| major | Last Consistent Snapshot of Slave peer '*local peer name*' could not be created. Maximal number of Volumes are already defined. | Delete Volumes to allow new ones to be created. Activate Mirror on the Master Machine. |

## MIRROR_END_SYNC_FAILED_CONFIGURATION_ERROR

| Severity | Description | Troubleshooting |
|---|---|---|
| major | Configuration of remote mirror of peer '*local peer name*' on target '*target name*' does not match local configuration. | Make sure configuration on both machines is compatible and activate the mirror. If problem persists contact IBM support. |

## MIRROR_CHANGE_DESIGNATION

| Severity | Description |
|---|---|
| informational | Local peer '*local peer name*' switched its designated role with peer on Target '*target name*'. It is now *designation*. |

## MIRROR_CANCEL_SNAPSHOT

| Severity | Description |
|---|---|
| informational | All mirrored snapshots which were created for Mirror of peer '*local peer name*' and were not yet synchronized will not be mirrored in the remote machine. |

## DM_DEFINE

| Severity | Description |
|---|---|
| informational | Data Migration was defined to Volume '*local volume name*' from Target '*target name*'. |

## DM_SYNC_STARTED

| Severity | Description |
|---|---|
| informational | Migration to volume '*local volume name*' from Target '*target name*' has started. |

## DM_SYNC_ENDED

| Severity | Description |
|---|---|
| informational | Migration to volume '*local volume name*' from target '*target name*' is complete. |

## DM_SYNC_ENDED_WITH_ERRORS

| Severity | Description |
|---|---|
| Critical | Migration to volume '*local volume name*' from target '*target name*' has completed with *medium_errors_in_data_migration* error(s). Check previous events related to this volume for the list of affected LBAs.'. |

## DM_ACTIVATE

| Severity | Description |
|---|---|
| informational | Migration to Volume '*local volume name*' from Target '*target name*' was activated. |

## DM_DEACTIVATE

| Severity | Description |
|---|---|
| informational | Migration to Volume '*local volume name*' from Target '*target name*' was deactivated. |

## DM_DEACTIVATE_LUN_UNAVAILABLE

| Severity | Description |
|---|---|
| minor | Migration to Volume '*local volume name*' from Target '*target name*' was deactivated since LUN is not available on one of the active paths to the target. |

## DM_DELETE

| Severity | Description |
|---|---|
| informational | Definition of Data Migration to Volume '*local volume name*' from Target '*target name*' was deleted. |

## SCHEDULE_CREATE

| Severity | Description |
|---|---|
| informational | Schedule was created with name '*schedule name*'. |

## SCHEDULE_UPDATE

| Severity | Description |
|---|---|
| informational | Schedule with name '*schedule name*' was updated. |

## SCHEDULE_RENAME

| Severity | Description |
|---|---|
| informational | Schedule with name '*old_name*' was renamed '*schedule name*'. |

## SCHEDULE_DELETE

| Severity | Description |
|---|---|
| informational | Schedule with name '*schedule name*' was deleted. |

## MIRROR_RPO_OK

| Severity | Description |
|---|---|
| informational | Mirror of local peer '*local peer name*' is now ahead of its specified RPO. |

## MIRROR_RPO_LAGGING

| Severity | Description |
|---|---|
| informational | Mirror of local peer '*local peer name*' is now behind its specified RPO. |

## MIRROR_CHANGE_RPO

| Severity | Description |
|---|---|
| informational | RPO or Mirror of local peer '*local peer name*' is now *RPO*. |

## MIRROR_IS_LAGGING_BEYOND_PERCENT_THRESHOLD

| Severity | Description |
|---|---|
| warning | Last Replication Time of Mirror of local peer '*local peer name*' is *Last Replication Time*. |

## MIRROR_AUTO_FIX_REACHED_LIMIT

| Severity | Description |
|---|---|
| warning | A remote checksum diff for mirror '*local peer name*' cannot be fixed automatically because we reached the auto fix limit. |

## MIRROR_IS_LAGGING_BEYOND_ABSOLUTE_THRESHOLD

| Severity | Description |
|---|---|
| warning | Last Replication Time of Mirror of local peer '*local peer name*' is *Last Replication Time*. |

## MIRROR_INCOMPATIBLE_VERSION_FOR_UNMAP_SUPPORT

| Severity | Description |
|---|---|
| warning | Mirror of peer '*local peer name*' on target '*target name*' cannot support unmap, remote machine has incompatible version. |

## MULTISITE_DEFINE

| Severity | Description |
|---|---|
| informational | A Multi-site Master was defined for volume '*local volume name*' on SMaster target and vol '*SMaster target name*' '*SMaster volume name*', and Slave target and vol '*slave target name*' '*slave volume name*'. |

## MULTISITE_DEFINE_SLAVE

| Severity | Description |
|---|---|
| informational | A Multi-site Slave was defined for volume '*local volume name*' on Master target and vol '*master target name*' '*master volume name*', and SMaster target and vol '*SMaster target name*' '*SMaster volume name*'. |

## MULTISITE_DEFINE_SMASTER

| Severity | Description |
|---|---|
| informational | A Multi-site SMaster was defined for volume '*local volume name*' on Master target and vol '*master target name*' '*master volume name*', and Slave target and vol '*slave target name*' '*slave volume name*'. |

## CG_MULTISITE_DEFINE

| Severity | Description |
|---|---|
| informational | A Multi-site Master was defined for consistency group '*local CG name*' on SMaster target and CG '*SMaster target name* '*SMaster CG name*', and Slave target and vol '*slave target name*' '*slave CG name*. |

## CG_MULTISITE_DEFINE_SLAVE

| Severity | Description |
|---|---|
| informational | A Multi-site Slave was defined for consistency group '*local CG name*' on Master target and CG '*master target name*' '*master CG name*', and SMaster target and CG '*SMaster target name*' '*SMaster CG name*'. |

## CG_MULTISITE_DEFINE_SMASTER

| Severity | Description |
|---|---|
| informational | A Multi-site SMaster was defined for consistency group '*local CG name*' on Master target and vol '*master target name*' '*master CG name*', and Slave target and CG '*slave target name*' '*slave CG name*'. |

## MULTISITE_DELETE

| Severity | Description |
|---|---|
| informational | The Multi-site relation of peer '*local peer name*' was deleted. |

## MULTISITE_CHANGE_SLAVE_ROLE_TO_MASTER

| Severity | Description |
|---|---|
| informational | The Multi-site relation of peer '*local peer name*' was changed to standalone. |

## MULTISITE_CHANGE_MASTER_ROLE_TO_SLAVE

| Severity | Description |
|---|---|
| informational | The Multi-site relation of peer '*local peer name*' was changed to Slave. |

## MULTISITE_CHANGE_SMASTER_ROLE_TO_MASTER

| Severity | Description |
|---|---|
| informational | The Multi-site relation of peer '*local peer name*' was changed to Master. |

## MULTISITE_CHANGE_MASTER_ROLE_TO_SMASTER

| Severity | Description |
|---|---|
| informational | The Multi-site relation of peer '*local peer name*' was changed to SMaster. |

## MULTISITE_STANDBY_MIRROR_REGISTERED

| Severity | Description |
|---|---|
| informational | The Multi-site relation of peer '*local peer name*' registered a standby mirror on SMaster system |

## MULTISITE_COMPROMISED

| Severity | Description |
|---|---|
| warning | The state of the Multi-site relation on peer '*local peer name*' changed to compromised, reason: *Compromise Reason* |

## MULTISITE_RESTORED

| Severity | Description |
|---|---|
| informational | The state of the Multi-site relation on peer '*local peer name*' changed from compromised to operational or degraded |

## MULTISITE_SWITCH_ROLES_TO_MASTER

| Severity | Description |
|---|---|
| informational | Local peer '*local peer name*' is now Master of Multi-site with SMaster target '*SMaster target name*', Slave target '*slave target name*'. |

## MULTISITE_SWITCH_ROLES_TO_SMASTER

| Severity | Description |
|---|---|
| informational | Local peer '*local peer name*' is now SMaster of Multi-site with Master target '*master target name*', Slave target '*slave target name*'. |

## MULTISITE_AUTOMATIC_ACTIVATION_OF_STANDBY_ASYNC_FAILED

| Severity | Description |
|---|---|
| major | Automatic activation of the standby Async mirror relation '*slave volume name*' had failed . |

## MAP_VOLUME

| Severity | Description |
|---|---|
| informational | Volume with name '*volume.name*' was mapped to LUN '*LUN*' for *host_or_cluster* with name '*host*'. |

## MAP_PROXY_VOLUME

| Severity | Description |
|---|---|
| informational | IBM Hyper-Scale Mobility Volume with name '*name*' was mapped to LUN '*LUN*' for *host_or_cluster* with name '*host*'. |

## UNMAP_VOLUME

| Severity | Description |
|---|---|
| informational | Volume with name '*volume.name*' was unmapped from *host_or_cluster* with name '*host*'. |

## UNMAP_PROXY_VOLUME

| Severity | Description |
|---|---|
| informational | IBM Hyper-Scale Mobility Volume with name '*name*' was unmapped from *host_or_cluster* with name '*host*'. |

## SPECIAL_TYPE_SET

| Severity | Description |
|---|---|
| informational | Type of *host_or_cluster* with name '*host*' was set to '*type*'. |

## SERVICE_HAS_FAILED

| Severity | Description | Troubleshooting |
|---|---|---|
| variable | *Component ID* has failed. | Contact IBM Support |

## SERVICE_FAILED_TO_PHASEIN

| Severity | Description | Troubleshooting |
|---|---|---|
| major | *Component ID* failed to phase-in. | Contact IBM Support |

## SERVICE_FAILED_TO_RESTART

| Severity | Description | Troubleshooting |
|---|---|---|
| major | *Component ID* failed to restart. | Contact IBM Support |

## MODULE_FAILED

| Severity | Description | Troubleshooting |
|---|---|---|
| critical | *Component ID* failed. | Contact IBM Support |

## NODE_OK

| Severity | Description |
|---|---|
| informational | Node #*Node ID* of type *Node Type* on *Module Component ID* has been revived and OK now. The node pid: *Process ID* |

## NODE_AUTO_RESET_FAILED

| Severity | Description | Troubleshooting |
|---|---|---|
| minor | Node #*Node ID* of type *Node Type* on *Module Component ID* has failed and failed to auto-revive. | Contact IBM Support |

## DISK_HAS_FAILED

| Severity | Description | Troubleshooting |
|---|---|---|
| variable | Disk *Component ID* Failed. | Please contact your Administrator. |

### SSD_HAS_FAILED

| Severity | Description | Troubleshooting |
|---|---|---|
| major | SSD *Component ID* Failed. | Please contact your Administrator. |

### VAULT_DEVICE_HAS_FAILED

| Severity | Description | Troubleshooting |
|---|---|---|
| minor | Vault device *Component ID* Failed. | Please contact your Administrator. |

### COMPONENT_TEST_OF_DISK_HAS_FAILED

| Severity | Description | Troubleshooting |
|---|---|---|
| major | Test of *Component ID* has failed with error *Error*. | Please contact your Administrator. |

### COMPONENT_TEST_OF_SSD_HAS_FAILED

| Severity | Description | Troubleshooting |
|---|---|---|
| major | Test of *Component ID* has failed with error *Error*. | Please contact your Administrator. |

### COMPONENT_TEST_OF_BOOT_MEDIA_HAS_FAILED

| Severity | Description | Troubleshooting |
|---|---|---|
| major | Test of *Component ID* has failed with error *Error*. | Please contact your Administrator. |

### COMPONENT_TEST_OF_VAULT_DEVICE_HAS_FAILED

| Severity | Description | Troubleshooting |
|---|---|---|
| major | Test of *Component ID* has failed with error *Error*. | Please contact your Administrator. |

### BOOT_MEDIA_COMPONENT_TEST_FAILED

| Severity | Description | Troubleshooting |
|---|---|---|
| major | Test of *Component ID* has failed with error *Error*. | Please contact your Administrator. |

### DISK_STARTED_PHASEOUT

| Severity | Description |
|---|---|
| informational | System started phasing out *Component ID*. |

### DISK_STARTED_PHASEIN

| Severity | Description |
|---|---|
| informational | System started phasing in *Component ID*. |

## DISK_FINISHED_PHASEIN

| Severity | Description |
|---|---|
| informational | System finished phasing in *Component ID*. |

## DISK_FINISHED_PHASEOUT

| Severity | Description |
|---|---|
| informational | System finished phasing out *Component ID*. |

## DISK_RECOVERED

| Severity | Description |
|---|---|
| critical | Disk *Component ID* is functioning again. |

## MODULE_STARTED_PHASEOUT

| Severity | Description |
|---|---|
| informational | System started phasing out *Component ID*. |

## MODULE_FINISHED_PHASEOUT

| Severity | Description |
|---|---|
| informational | System finished phasing out *Component ID*. |

## MODULE_STOPPED_PHASEOUT_DUE_TO_MANAGEMENT_REQUIREMENT

| Severity | Description |
|---|---|
| major | System stopped phasing out *Component ID* due to management requirement. |

## START_WORK

| Severity | Description |
|---|---|
| informational | System has entered ON state. |

## SYSTEM_HAS_ENTERED_MAINTENANCE_MODE

| Severity | Description |
|---|---|
| warning | System has entered MAINTENANCE state [*Reason*] |

## SYSTEM_LEFT_CHARGING_STATE

| Severity | Description |
|---|---|
| informational | System is sufficiently charged. |

## USER_SHUTDOWN

| Severity | Description |
|---|---|
| major | System is shutting down due to a user request. |

## EMERGENCY_SHUTDOWN_NOW

| Severity | Description | Troubleshooting |
|---|---|---|
| critical | System is shutting down in emergency shutdown mode due to: *Emergency Shutdown Reason*. | Please contact your Administrator. |

## SHOULD_HAVE_BEEN_EMERGENCY_SHUTDOWN

| Severity | Description |
|---|---|
| critical | An emergency shutdown has been detected, but emergency shutdown is disabled for the detected reason. Shutdown reason: *Shutdown Reason*. |

## DATA_SERVICE_STARTED_PHASEOUT

| Severity | Description |
|---|---|
| informational | System started phasing out *Component ID*. |

## DATA_SERVICE_STARTED_PHASEIN

| Severity | Description |
|---|---|
| informational | System started phasing in *Component ID*. |

## DATA_SERVICE_FINISHED_PHASEIN

| Severity | Description |
|---|---|
| informational | System finished phasing in *Component ID*. |

## DATA_SERVICE_FINISHED_PHASEOUT

| Severity | Description |
|---|---|
| informational | System finished phasing out *Component ID*. |

## TXN_REBUILD_STARTED

| Severity | Description |
|---|---|
| informational | Start rebuild process of txns. |

## TXN_REBUILD_ENDED

| Severity | Description |
|---|---|
| informational | End rebuild process for txns. |

## TXN_REDIST_STARTED

| Severity | Description |
|---|---|
| informational | Start redist process for txns. |

### TXN_REDIST_ENDED

| Severity | Description |
|---|---|
| informational | End redist process for txns. |

### DISK_MARKED_TO_PHASEOUT

| Severity | Description |
|---|---|
| informational | System started phasing out *Component ID*. |

### DISK_MARKED_TO_PHASEIN

| Severity | Description |
|---|---|
| informational | System started phasing out *Component ID*. |

### CANNOT_CREATE_NEW_DATA_DISTRIBUTION

| Severity | Description |
|---|---|
| informational | System cannot phaseout disks for storage medium *Storage Medium*. |

### CANNOT_RESIZE_FLASH_MEDIUM_POOLS

| Severity | Description |
|---|---|
| informational | System failed to resize flash medium pools for TMS phaseout. |

### DATA_PROTECTION_STATUS_CHANGED

| Severity | Description |
|---|---|
| variable | Data protection status has changed from '*old_status*' to '*new_status*' |

### VAULT_DEVICE_SECURE_ERASE_PROCESS_SUCCESSFUL

| Severity | Description |
|---|---|
| informational | Vault device secure erase process successful. |

### VAULT_DEVICE_SECURE_ERASE_PROCESS_FAILED

| Severity | Description |
|---|---|
| major | Vault device secure erase process failed [*Reason*] |

### VAULT_DEVICE_FAILED_SECURE_ERASE

| Severity | Description |
|---|---|
| major | Secure erase for *Component ID* failed. [*Reason*]. |

### SYSTEM_PHYSICAL_CAPACITY_CHANGED

| Severity | Description |
|---|---|
| informational | System physical capacity is now *Capacity*GB. |

## SYSTEM_EFFECTIVE_CAPACITY_CHANGED

| Severity | Description |
|---|---|
| informational | System effective capacity is now *Capacity*GB. |

## SYSTEM_OUT_OF_PHYSICAL_SPACE

| Severity | Description |
|---|---|
| critical | System has run out of physical capacity. All volumes are now write-protected. |

## SYSTEM_NORMAL_OPERATION_RESUMED

| Severity | Description |
|---|---|
| informational | Normal operation is resumed. Volumes have been restored to their original write-protection state. |

## SYSTEM_FREE_PHYSICAL_CAPACITY_REACHED_RESERVED_LIMIT

| Severity | Description |
|---|---|
| critical | System physical capacity usage reached reserved limit. All volumes are now write-protected. |

## ENCRYPT_ENABLE_STARTED

| Severity | Description |
|---|---|
| informational | Starting encryption activation. This process can take several minutes to complete. |

## ENCRYPT_ENABLE_COMPLETED

| Severity | Description |
|---|---|
| informational | Encryption is in effect. |

## ENCRYPT_ENABLE_NOT_COMPLETED

| Severity | Description | Troubleshooting |
|---|---|---|
| major | Cannot complete encryption activation because *reason*. *Count* vault device(s) and *Count* flash enclosure(s) could not be enrolled. | Please contact technical support |

## ENCRYPT_DISABLE_STARTED

| Severity | Description |
|---|---|
| informational | Starting encryption deactivation. This process can take several minutes to complete. |

## ENCRYPT_DISABLE_COMPLETED

| Severity | Description |
|---|---|
| informational | Encryption is no longer in effect. |

## ENCRYPT_DISABLE_NOT_COMPLETED

| Severity | Description | Troubleshooting |
|---|---|---|
| major | Cannot complete encryption deactivation because *reason*. *Count* vault device(s) and *Count* flash enclosure(s) could not be crypto erased. | Please contact technical support |

## ENCRYPT_KEYSERVER_ADDED

| Severity | Description |
|---|---|
| informational | A key server named '*Key Server Name*' was added. |

## ENCRYPT_KEYSERVER_DELETED

| Severity | Description |
|---|---|
| informational | Key server '*Key Server Name*' was deleted. |

## ENCRYPT_KEYSERVER_EDITED

| Severity | Description |
|---|---|
| informational | Details of key server '*Key Server Name*' were modified. |

## ENCRYPT_KEYSERVER_RENAMED

| Severity | Description |
|---|---|
| informational | Key server '*Old Name*' was renamed to '*New Name*'. |

## ENCRYPT_KEYSERVER_CHECK_STATUS_STARTED

| Severity | Description |
|---|---|
| informational | Start checking connectivity status of the keyservers currently defined in the system. This process can take several minutes to complete. |

## ENCRYPT_KEYSERVER_CHECK_STATUS_COMPLETED

| Severity | Description |
|---|---|
| informational | Completed checking connectivity status of the keyservers currently defined in the system. |

## ENCRYPT_KEYSERVER_REKEY_COMPLETED

| Severity | Description |
|---|---|
| informational | Key server '*Key Server Name*' rekey completed. |

## ENCRYPT_LOCAL_REKEY_COMPLETED

| Severity | Description |
|---|---|
| informational | Local key rekey completed. |

### ENCRYPT_CHANGE_KEY_SCHEME_COMPLETED

| Severity | Description |
|---|---|
| informational | Change key scheme from external to local key completed. |

### ENCRYPT_CHANGE_KEY_SCHEME_FAILED

| Severity | Description |
|---|---|
| major | Change key scheme from external to local key failed because *failure reason*. |

### ENCRYPT_CHANGE_KEY_SCHEME_ROLLBACK_FAILED

| Severity | Description |
|---|---|
| major | Cannot rollback change key scheme. |

### ENCRYPT_UNABLE_TO_UPDATE_KEY_DURING_DEACTIVATE_ON_KEYSERVER

| Severity | Description |
|---|---|
| major | Could not update key server '*Keyserver Name*' regarding encryption deactivation. Please check key server status. |

### ENCRYPT_KEYSERVER_REKEY_FAILED

| Severity | Description |
|---|---|
| major | Cannot complete rekey with key server '*Key Server Name*'. |

### ENCRYPT_LOCAL_REKEY_FAILED

| Severity | Description |
|---|---|
| major | Local key rekey failed because *failure reason*. |

### ENCRYPT_KEYSERVER_REKEY_ROLLBACK_FAILED

| Severity | Description |
|---|---|
| major | Cannot rollback failed rekey with key server '*Key Server Name*'. |

### ENCRYPT_LOCAL_REKEY_ROLLBACK_FAILED

| Severity | Description |
|---|---|
| major | Cannot rollback failed local key rekey. |

### ENCRYPT_RECOVERY_KEY_ENTERED

| Severity | Description |
|---|---|
| informational | Valid recovery key share was entered by user '*User Name*'. |

### ENCRYPT_INVALID_RECOVERY_KEY_ENTERED

| Severity | Description |
|---|---|
| major | Invalid recovery key share was entered by user '*User Name*'. |

## ENCRYPT_RECOVERY_KEYS_GENERATED

| Severity | Description |
|---|---|
| informational | Recovery keys created. |

## ENCRYPT_RECOVERY_KEY_REKEY_SUCCESS

| Severity | Description |
|---|---|
| informational | Recovery key rekey was successful. |

## ENCRYPT_RECOVERY_KEY_REKEY_FAIL

| Severity | Description |
|---|---|
| major | Recovery key rekey failed. |

## ENCRYPT_RECOVERY_KEY_VERIFIED

| Severity | Description |
|---|---|
| informational | Recovery key verified successfully for user '*User Name*'. |

## ENCRYPT_RECOVERY_KEY_VERIFY_FAILED

| Severity | Description |
|---|---|
| major | Recovery key verification failed for user '*User Name*'. |

## ENCRYPT_RECOVERY_KEY_ALL_SHARES_VERIFIED

| Severity | Description |
|---|---|
| informational | All recovery key shares have been verified. |

## ENCRYPT_KR_WRITE_FAILED

| Severity | Description | Troubleshooting |
|---|---|---|
| critical | Key repository write failed with error code: *rc*. | Please contact technical support |

## ENCRYPT_KR_READ_FAILED

| Severity | Description | Troubleshooting |
|---|---|---|
| major | Key repository read failed with error code: *rc*. | Please contact technical support |

## ENCRYPT_UNABLE_TO_RETRIEVE_KEY_FROM_KEYSERVER

| Severity | Description |
|---|---|
| major | Failed to retrieve key from key server '*Keyserver Name*' via *TEXT* on module *node id*. Please verify that the key server type and version are supported. If so, please check its status. |

## ENCRYPT_RECOVERY_KEY_RECOVER_SUCCESSFUL

| Severity | Description |
|---|---|
| informational | Key recovery was successful, unlocking system. |

## ENCRYPTION_CERTIFICATE_FOR_XIV_IS_NOT_INSTALLED

| Severity | Description | Troubleshooting |
|---|---|---|
| critical | XIV certificate is not installed. | Check output of pki_list for a certificate named XIV and contact technical support |

## ENCRYPT_UNABLE_TO_DELETE_MASTER_KEYSERVER

| Severity | Description |
|---|---|
| informational | Deletion of master key server '*Keyserver Name*' is not allowed. Please define another key server as master first'. |

## ENCRYPTION_SKMIP_ERROR

| Severity | Description | Troubleshooting |
|---|---|---|
| major | Module *Module* reported *Keyserver Name* returned error:*error code - TEXT* | Please contact the next level of support. |

## WIPEOUT_STARTED

| Severity | Description |
|---|---|
| informational | Starting the wipeout process. This process may take several minutes to complete. |

## WIPEOUT_COMPLETED

| Severity | Description |
|---|---|
| informational | The wipeout process finished successfully. |

## WIPEOUT_NOT_COMPLETED

| Severity | Description | Troubleshooting |
|---|---|---|
| major | Cannot complete the wipeout process because *reason*. *Count* vault device(s) and *Count* flash enclosure(s) could not be crypto erased. | Contact IBM Support |

## DIMM_FAILED

| Severity | Description | Troubleshooting |
|---|---|---|
| major | *Component ID* has failed. Hardware status: *Status*. | Contact IBM Support |

## CPU_FAILED

| Severity | Description | Troubleshooting |
|---|---|---|
| major | *Component ID* has failed. Hardware status: *Status*. | Contact IBM Support |

## NIC_FAILED

| Severity | Description | Troubleshooting |
|---|---|---|
| major | *Component ID* has failed. Hardware status: *Status*. | Contact IBM Support |

## MODULE_BBU_FAILED

| Severity | Description | Troubleshooting |
|---|---|---|
| major | *BBU id* has failed. Hardware status: '*Status*'. BBU state: '*State*'. | Contact IBM Support |

## MODULE_BBU_OK

| Severity | Description |
|---|---|
| informational | *BBU id* is now OK. |

## DIMM_WAS_REMOVED

| Severity | Description | Troubleshooting |
|---|---|---|
| minor | The DIMM with serial number '*Serial*' was removed from *ModuleId*. | Was this DIMM actually removed? |

## PSU_CHANGE_DETECTED

| Severity | Description | Troubleshooting |
|---|---|---|
| informational | *Component ID* has been changed from a serial number '*old_serial*', part number '*old_part_number*', to serial number '*new_serial*' and part number '*new_part_number*'. | Was this PSU actually replaced? |

## PSU_WAS_REMOVED

| Severity | Description | Troubleshooting |
|---|---|---|
| warning | *Component ID* with a serial number '*Serial*' and part number '*Part Number*' was removed from the system. | Was this PSU actually removed? |

## PSU_MISSING_FROM_INSTALL

| Severity | Description | Troubleshooting |
|---|---|---|
| major | System was installed without *Component ID*. | This PSU was not in the system at install time. |

## PSU_WAS_INSTALLED

| Severity | Description | Troubleshooting |
|---|---|---|
| informational | *Component ID* with a serial number '*Serial*' and part number '*Part Number*' was installed in the system. | Was this PSU actually installed? |

## NIC_CHANGE_DETECTED

| Severity | Description | Troubleshooting |
|---|---|---|
| major | *Component ID* has been changed from a serial of *old_serial* to *new_serial*. | Was this NIC actually replaced? |

## MODULE_BBU_IS_TOO_OLD

| Severity | Description |
|---|---|
| major | *BBU id* installed more than *max use time* months after manufacturing date *manufacturing date*. |

## MODULE_BBU_SHELF_TIME_EXCEEDED

| Severity | Description |
|---|---|
| major | *BBU id* installed more than *max shelf time* months after manufacturing date *manufacturing date*. |

## TECHNICIAN_WORK_STARTED

| Severity | Description |
|---|---|
| informational | Technician work has started, expected to end at *End Time*. Comment: *Comment*. |

## TECHNICIAN_WORK_ENDED

| Severity | Description |
|---|---|
| informational | Technician work has ended after *Elapsed Time* minutes. Comment: *Comment*. |

## TECHNICIAN_WORK_TIMED_OUT

| Severity | Description |
|---|---|
| warning | Technician work has timed out after *Elapsed Time* minutes. Comment: *Comment*. |

## XIV_SUPPORT_ENABLED

| Severity | Description |
|---|---|
| informational | XIV support access from *From* is enabled from *Start Time* until *Finish Time*. Comment: *Comment*. |

## XIV_SUPPORT_ENABLED_NO_TIME_LIMIT

| Severity | Description |
|---|---|
| informational | XIV support access from *From* is enabled from *Start Time* until explicitly disabled. Comment: *Comment*. |

## XIV_SUPPORT_DISABLED

| Severity | Description |
|---|---|
| informational | XIV support access is disabled. |

## XIV_SUPPORT_WINDOW_TIMEOUT

| Severity | Description |
|---|---|
| informational | XIV support work window timeout is expired. |

## HOST_DEFINE

| Severity | Description |
|---|---|
| informational | Host of type *host.type* was defined with name '*host.name*'. |

## HOST_UPDATE

| Severity | Description |
|---|---|
| informational | Host named '*host.name*' was updated. |

## CLUSTER_CREATE

| Severity | Description |
|---|---|
| informational | Cluster was defined with name '*cluster.name*'. |

## HOST_DEFINE_FAILED_TOO_MANY

| Severity | Description | Troubleshooting |
|---|---|---|
| warning | Host with name '*name*' could not be defined. You are attempting to define more hosts than the system permits. | Delete Hosts to allow new ones to be defined. |

## CLUSTER_CREATE_FAILED_TOO_MANY

| Severity | Description | Troubleshooting |
|---|---|---|
| warning | Cluster with name '*name*' could not be defined. You are attempting to define more Clusters than the system permits. | Delete Clusters to allow new ones to be defined. |

## HOST_RENAME

| Severity | Description |
|---|---|
| informational | Host with name '*old_name*' was renamed '*host.name*'. |

## CLUSTER_RENAME

| Severity | Description |
|---|---|
| informational | Cluster with name '*old_name*' was renamed '*cluster.name*'. |

## HOST_DELETE

| Severity | Description |
|---|---|
| informational | Host with name '*host.name*' was deleted. |

## CLUSTER_DELETE

| Severity | Description |
|---|---|
| informational | Cluster with name '*cluster.name*' was deleted. |

## HOST_ADD_PORT

| Severity | Description |
|---|---|
| informational | Port of type *type* and ID '*port_name*' was added to Host with name '*host.name*'. |

## CLUSTER_ADD_HOST

| Severity | Description |
|---|---|
| informational | Host with name '*host.name*' was added to Cluster with name '*cluster.name*'. |

## HOST_REMOVE_PORT

| Severity | Description |
|---|---|
| informational | Port of type *type* and ID '*port_name*' was removed from Host with name '*host.name*' was deleted. |

## CLUSTER_REMOVE_HOST

| Severity | Description |
|---|---|
| informational | Host with name '*host.name*' was removed from Cluster with name '*cluster.name*'. |

## DESTINATION_DEFINE

| Severity | Description |
|---|---|
| informational | Destination with name '*name*' was defined. |

## DESTINATION_UPDATE

| Severity | Description |
|---|---|
| informational | Destination with name '*name*' was updated. |

## DESTINATION_DELETE

| Severity | Description |
|---|---|
| informational | Destination with name '*name*' was deleted. |

## DESTINATION_RENAME

| Severity | Description |
|---|---|
| informational | Destination with name '*old name*' was renamed '*new name*'. |

## DESTINATION_GROUP_CREATE

| Severity | Description |
|---|---|
| informational | Destination Group with name '*name*' was created. |

## DESTINATION_GROUP_UPDATE

| Severity | Description |
|---|---|
| informational | Destination Group with name '*name*' was updated. |

## DESTINATION_GROUP_DELETE

| Severity | Description |
|---|---|
| informational | Destination Group with name '*name*' was deleted. |

## DESTINATION_GROUP_RENAME

| Severity | Description |
|---|---|
| informational | Destination Group with name '*old name*' was renamed '*new name*'. |

## DESTINATION_GROUP_ADD_DESTINATION

| Severity | Description |
|---|---|
| informational | Destination with name '*destination name*' was added to destination group '*destgroup name*'. |

## DESTINATION_GROUP_REMOVE_DESTINATION

| Severity | Description |
|---|---|
| informational | Destination with name '*destination name*' was removed from destination group '*destgroup name*'. |

## RULE_CREATE

| Severity | Description |
|---|---|
| informational | Rule with name '*name*' was created. |

## RULE_UPDATE

| Severity | Description |
|---|---|
| informational | Rule with name '*name*' was updated. |

## RULE_DELETE

| Severity | Description |
|---|---|
| informational | Rule with name '*name*' was deleted. |

## RULE_RENAME

| Severity | Description |
|---|---|
| informational | Rule with name '*old name*' was renamed '*new name*'. |

## SMTP_GATEWAY_DEFINE

| Severity | Description |
|---|---|
| informational | SMTP gateway with name '*name*' was defined. |

## SMTP_GATEWAY_UPDATE

| Severity | Description |
|---|---|
| informational | SMTP gateway with name '*name*' was updated. |

## SMTP_GATEWAY_DELETE

| Severity | Description |
|---|---|
| informational | SMTP gateway with name '*name*' was deleted. |

## SMTP_GATEWAY_RENAME

| Severity | Description |
|---|---|
| informational | SMTP gateway with name '*old name*' was renamed '*new name*'. |

## SMTP_GATEWAY_PRIORITIZE

| Severity | Description |
|---|---|
| informational | SMTP gateways were prioritized; the new order is *order*. |

## CALL_HOME_CONNECTION_OK

| Severity | Description |
|---|---|
| informational | Events are sent to the Call Home server by SMTP gateway '*name*'. |

## CALL_HOME_CONNECTION_PROBLEM

| Severity | Description | Troubleshooting |
|---|---|---|
| major | Events are not sent to the Call Home server by SMTP gateway '*name*'. Reason: *Event Reason*. | Please contact IBM support. |

## SMTP_GATEWAY_FAILED

| Severity | Description |
|---|---|
| major | SMTP gateway with name '*name*' has failed. It will not be used until *Retry Time*. |

## SMTP_GATEWAY_VIA_NODE_FAILED

| Severity | Description |
|---|---|
| warning | Sending event *Event Code* (*Event Index*) through *SMTP Gateway* via *Module ID* has failed; Error message: '*Error Message*'. |

## SMS_GATEWAY_DEFINE

| Severity | Description |
|---|---|
| informational | SMS gateway with name '*name*' was defined. |

## SMS_GATEWAY_UPDATE

| Severity | Description |
|---|---|
| informational | SMS gateway with name '*name*' was updated. |

## SMS_GATEWAY_DELETE

| Severity | Description |
|---|---|
| informational | SMS gateway with name '*name*' was deleted. |

## SMS_GATEWAY_RENAME

| Severity | Description |
|---|---|
| informational | SMS gateway with name '*old name*' was renamed '*new name*'. |

## SMS_GATEWAY_PRIORITIZE

| Severity | Description |
|---|---|
| informational | SMS gateways were prioritized; the new order is *order*. |

## CONS_GROUP_CREATE

| Severity | Description |
|---|---|
| informational | Consistency Group with name '*cg.name*' was created. |

## CONS_GROUP_CREATE_FAILED_TOO_MANY

| Severity | Description | Troubleshooting |
|---|---|---|
| warning | Consistency Group with name '*name*' could not be created. You are attempting to add more Consistency Groups than the system permits. | Delete Consistency Groups to allow new ones to be created. |

## CONS_GROUP_RENAME

| Severity | Description |
|---|---|
| informational | Consistency Group with name '*old_name*' was renamed '*cg.name*'. |

## SECONDARY_CONS_GROUP_RENAME

| Severity | Description |
|---|---|
| informational | Consistency Group with name '*old_name*' was renamed '*cg.name*' by primary machine. |

## CONS_GROUP_DELETE

| Severity | Description |
|---|---|
| informational | Consistency Group with name '*cg.name*' was deleted. |

## CONS_GROUP_ADD_VOLUME

| Severity | Description |
|---|---|
| informational | Volume with name '*volume.name*' was added to Consistency Group with name '*cg.name*'. |

## SLAVE_CONS_GROUP_ADD_VOLUME

| Severity | Description |
|---|---|
| informational | Volume with name '*volume.name*' was added to Consistency Group with name '*cg.name*' by its remote peer. |

## CONS_GROUP_REMOVE_VOLUME

| Severity | Description |
|---|---|
| informational | Volume with name '*volume.name*' was removed from Consistency Group with name '*cg.name*'. |

## SLAVE_CONS_GROUP_REMOVE_VOLUME

| Severity | Description |
|---|---|
| informational | Volume with name '*volume.name*' was removed from Consistency Group with name '*cg.name*' by its remote peer. |

## CONS_GROUP_SNAPSHOTS_CREATE

| Severity | Description |
|---|---|
| informational | Snapshot Group for Consistency Group with name '*cg.name*' was created with name '*sg.name*'. |

## CONS_GROUP_SNAPSHOTS_CREATE_FAILED_TOO_MANY

| Severity | Description | Troubleshooting |
|---|---|---|
| warning | Snapshot Group for Consistency Group '*cg.name*' could not be created. You are attempting to add more Volumes than the system permits. | Delete Volumes to allow new ones to be created. |

## CONS_GROUP_SNAPSHOTS_OVERWRITE

| Severity | Description |
|---|---|
| informational | Snapshot Group named '*sg.name*' was overriden for Consistency Group with name '*cg.name*'. |

## SLAVE_CONS_GROUP_SNAPSHOTS_CREATE

| Severity | Description |
|---|---|
| informational | Mirrored Snapshot Group for Consistency Group with name '*cg.name*' was created with name '*sg.name*'. |

## HA_SLAVE_CONS_GROUP_SNAPSHOTS_CREATE

| Severity | Description |
|---|---|
| informational | HyperSwap Snapshot Group for Consistency Group with name '*cg.name*' was created with name '*sg.name*'. |

## SLAVE_CONS_GROUP_SNAPSHOTS_OVERWRITE

| Severity | Description |
|---|---|
| informational | Mirrored Snapshot Group named '*sg.name*' was overriden for Consistency Group with name '*cg.name*'. |

## HA_SLAVE_CONS_GROUP_SNAPSHOTS_OVERWRITE

| Severity | Description |
|---|---|
| informational | HyperSwap Snapshot Group named '*sg.name*' was overriden for Consistency Group with name '*cg.name*'. |

## MIRROR_CONS_GROUP_SNAPSHOTS_CREATE

| Severity | Description |
|---|---|
| informational | Mirrored Snapshot Group for Consistency Group with name '*cg.name*' was created with name '*sg.name*'. |

## HA_CONS_GROUP_SNAPSHOTS_CREATE

| Severity | Description |
|---|---|
| informational | HyperSwap Snapshot Group for Consistency Group with name '*cg.name*' was created with name '*sg.name*'. |

## MIRROR_CONS_GROUP_SNAPSHOTS_OVERWRITE

| Severity | Description |
|---|---|
| informational | Mirrored Snapshot Group named '*sg.name*' was overriden for Consistency Group with name '*cg.name*'. |

## HA_CONS_GROUP_SNAPSHOTS_OVERWRITE

| Severity | Description |
|---|---|
| informational | HyperSwap Snapshot Group named '*sg.name*' was overriden for Consistency Group with name '*cg.name*'. |

## REMOTE_MIRROR_CONS_GROUP_SNAPSHOTS_NOT_CREATED_YET

| Severity | Description |
|---|---|
| minor | Remote Mirrored Snapshot Group for Consistency Group with name '*remote CG name*' with name '*cg_sync_job.remote_snapgroup*' on Target '*target name*' were not created yet. |

## MIRROR_SNAPGROUP_CREATE_FAILED

| Severity | Description |
|---|---|
| minor | Remote snapshot group named '*snapshot group name*' was not created successfully. Error code is '*error*' |

## SNAPSHOT_GROUP_RESTORE

| Severity | Description |
|---|---|
| informational | Volumes were restored from Snapshot Group with name '*sg.name*'. |

## SNAPSHOT_GROUP_RENAME

| Severity | Description |
|---|---|
| informational | Snapshot Group with name '*old_sg.name*' were renamed to '*sg.name*'. |

## SNAPSHOT_GROUP_DUPLICATE

| Severity | Description |
|---|---|
| informational | All Snapshots in Snapshot Group with name '*sg.name*' were duplicated. Duplicate Snapshot Group is named '*sg.name*'. |

## SNAPSHOT_GROUP_FORMAT

| Severity | Description |
|---|---|
| informational | All Snapshots in Snapshot Group with name '*sg.name*' were formatted'. |

## SNAPSHOT_GROUP_DELETE

| Severity | Description |
|---|---|
| informational | All Snapshots in Snapshot Group with name '*sg.name*' were deleted. |

## SNAPSHOT_GROUP_CHANGE_PRIORITY

| Severity | Description |
|---|---|
| informational | Deletion Priority of all Snapshots in Snapshot Group with name '*sg.name*' were changed from '*old priority*' to '*new priority*'. |

## SNAPSHOT_GROUP_LOCK

| Severity | Description |
|---|---|
| informational | All Snapshots in Snapshot Group with name '*sg.name*' were locked. |

## SNAPSHOT_GROUP_UNLOCK

| Severity | Description |
|---|---|
| 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. |

## TRANSACTION_NODE_DOES_NOT_USE_OPTIMAL_SRP_PATH

| Severity | Description | Troubleshooting |
|---|---|---|
| warning | Data service '*service*' does not use the optimal path to '*enclosure*'. | Contact IBM Support |

## TRANSACTION_NODE_USES_OPTIMAL_SRP_PATH

| Severity | Description | Troubleshooting |
|---|---|---|
| informational | Data service '*service*' uses the optimal path to '*enclosure*'. | Contact IBM Support |

## USER_DEFINED

| Severity | Description |
|---|---|
| informational | A user with name '*Name*' and category *Category* was defined. |

## USER_DELETED

| Severity | Description |
|---|---|
| informational | A user with name '*Name*' and category *Category* was deleted. |

## USER_RENAMED

| Severity | Description |
|---|---|
| informational | User with name '*Old Name*' was renamed '*New Name*'. |

## USER_UPDATED

| Severity | Description |
|---|---|
| informational | User with name '*Name*' was updated. |

## USER_ADDED_TO_USER_GROUP

| Severity | Description |
|---|---|
| informational | User '*User Name*' was added to user group '*User Group Name*'. |

## USER_REMOVED_FROM_USER_GROUP

| Severity | Description |
|---|---|
| informational | User '*User Name*' was removed from user group '*User Group Name*'. |

## USER_GROUP_CREATED

| Severity | Description |
|---|---|
| informational | A user group with name '*Name*' was created. |

## USER_GROUP_DELETED

| Severity | Description |
|---|---|
| informational | A user group with name '*Name*' was deleted. |

## USER_GROUP_RENAMED

| Severity | Description |
|---|---|
| informational | User group with name '*Old Name*' was renamed '*New Name*'. |

## LDAP_AUTHENTICATION_ACTIVATED

| Severity | Description |
|---|---|
| informational | LDAP authentication activated. |

## LDAP_AUTHENTICATION_DEACTIVATED

| Severity | Description |
|---|---|
| warning | LDAP authentication deactivated. |

## LDAP_CONFIGURATION_CHANGED

| Severity | Description |
|---|---|
| warning | LDAP configuration has changed. |

## LDAP_CONFIGURATION_RESET

| Severity | Description |
|---|---|
| warning | LDAP configuration has reset. |

## USER_LOGIN_HAS_SUCCEEDED

| Severity | Description |
|---|---|
| informational | User '*User Name*' from IP '*Client Address*' successfully logged into the system. |

## USER_LOGIN_HAS_FAILED

| Severity | Description |
|---|---|
| warning | User '*User Name*' from IP '*Client Address*' failed logging into the system. |

## USER_HAS_FAILED_TO_RUN_COMMAND

| Severity | Description |
|---|---|
| warning | User '*User Name*' from IP '*Client Address*' failed authentication when trying to run command '*Command Line*'. |

## LDAP_SERVER_INACCESSIBLE

| Severity | Description |
|---|---|
| minor | LDAP server *FQDN* is inaccessible. |

## LDAP_SERVER_ACCESSIBLE

| Severity | Description |
|---|---|
| informational | LDAP server *FQDN* is now accessible. |

## LDAP_SSL_CERTIFICATE_ABOUT_TO_EXPIRE

| Severity | Description |
|---|---|
| variable | SSL Certificate of LDAP server '*Server FQDN*' is about to expire on *Expiration Date* (*Counter* notification). |

## LDAP_SERVER_WAS_ADDED

| Severity | Description |
|---|---|
| informational | LDAP server '*Server FQDN*' was added to the system. |

## LDAP_SERVER_WAS_REMOVED

| Severity | Description |
|---|---|
| informational | LDAP server '*Server FQDN*' was removed from the system. |

## DESIGNATED_MSM_USER

| Severity | Description |
|---|---|
| informational | *Description* |

## DOMAIN_POLICY_SET

| Severity | Description |
|---|---|
| informational | Domain policy for *Parameter Name* set to '*Parameter Value*' |

## USER_ADDED_TO_DOMAIN

| Severity | Description |
|---|---|
| informational | User *User Name* was added to domain *Domain Name* (*Exclusive*). |

## USER_REMOVED_FROM_DOMAIN

| Severity | Description |
|---|---|
| informational | User *User Name* was removed from domain *Domain Name*. |

## APPADMIN_CAPABILITIES_SET

| Severity | Description |
|---|---|
| informational | Application admin capabilities have been set to *Capabilities* |

## ACCESS_TO_HOST_GRANTED_TO_USER_GROUP

| Severity | Description |
|---|---|
| informational | User group '*User Group Name*' was granted access to host '*Host Name*'. |

## ACCESS_OF_USER_GROUP_TO_HOST_REMOVED

| Severity | Description |
|---|---|
| informational | Access of User group '*User Group Name*' to host '*Host Name*' was removed. |

## ACCESS_TO_CLUSTER_GRANTED_TO_USER_GROUP

| Severity | Description |
|---|---|
| informational | User group '*User Group Name*' was granted access to cluster '*Cluster Name*'. |

## ACCESS_OF_USER_GROUP_TO_CLUSTER_REMOVED

| Severity | Description |
|---|---|
| informational | Access of User group '*User Group Name*' to cluster '*Cluster Name*' was removed. |

## COMPONENT_TEST_HAS_FAILED

| Severity | Description | Troubleshooting |
|---|---|---|
| variable | Test of *Component ID* has failed. Failure reason: *Failure Reason*. | Contact IBM Support |

## COMPONENT_TEST_SUCCEEDED

| Severity | Description | Troubleshooting |
|---|---|---|
| informational | Test of *Component ID* succeeded. | Contact IBM Support |

## MODULE_COMPONENT_TEST_STARTED

| Severity | Description |
|---|---|
| informational | Test of *Component ID* started. |

## DISK_COMPONENT_TEST_STARTED

| Severity | Description |
|---|---|
| informational | Test of *Component ID* started. |

## IB_SWITCH_COMPONENT_TEST_STARTED

| Severity | Description |
|---|---|
| informational | Test of *Component ID* started. |

## SSD_COMPONENT_TEST_STARTED

| Severity | Description |
|---|---|
| informational | Test of *Component ID* started. |

## VAULT_DEVICE_COMPONENT_TEST_STARTED

| Severity | Description |
|---|---|
| informational | Test of *Component ID* started. |

## BOOT_MEDIA_COMPONENT_TEST_STARTED

| Severity | Description |
|---|---|
| informational | Test of *Component ID* started. |

## FLASH_COMPONENT_TEST_STARTED

| Severity | Description |
|---|---|
| informational | Test of *Component ID* started. |

## BOOT_MEDIA_FAILED

| Severity | Description | Troubleshooting |
|---|---|---|
| major | *Component ID* has failed. Hardware status: *Status*. | Contact IBM Support |

## MODULE_BBU_TEST_STARTED

| Severity | Description |
|---|---|
| informational | Test of *Component ID* started. |

## MODULE_BBU_TEST_WILL_CONTINUE

| Severity | Description |
|---|---|
| warning | Test of *Component ID* will continue. After waiting *Minutes* minutes, capacity is still *Capacity*%. |

## COMPONENT_WAS_PHASED_OUT

| Severity | Description |
|---|---|
| informational | *Component ID* was phased-out. |

## COMPONENT_WAS_FAILED

| Severity | Description |
|---|---|
| variable | Component *Component ID* was marked as failed. |

## COMPONENT_FAILURE_WAS_CANCELED

| Severity | Description |
|---|---|
| informational | Component *Component ID* failure status was reset. |

## COMPONENT_WAS_PHASED_IN

| Severity | Description |
|---|---|
| informational | *Component ID* was phased-in. |

## COMPONENT_WAS_EQUIPPED

| Severity | Description |
|---|---|
| informational | *Component ID* was equipped. |

## INTERFACE_SERVICES_ACTIVATED

| Severity | Description |
|---|---|
| informational | Interface services of *Module ID* were activated. |

## COMPONENT_FIRMWARE_UPGRADE_ABORTING

| Severity | Description |
|---|---|
| warning | Aborting *Upgrade type* upgrade of *Firmware type* firmware, version *Label*, on *Scope*. Abort reason: *Reason*. Waiting for current upgrade item to complete. |

## COMPONENT_FIRMWARE_UPGRADE_ABORTED

| Severity | Description |
|---|---|
| warning | Aborted *Upgrade type* upgrade of *Firmware type* firmware, version *Label*, on *Scope*. Abort reason: *Reason*. Progress *Attempted*/*Total*, *Successes* succeeded, *Failures* failed, *No-Ops* no-ops. |

## COMPONENT_FIRMWARE_UPGRADE_DONE

| Severity | Description |
|---|---|
| informational | Finished *Upgrade type* upgrade of *Firmware type* firmware, version *Label*, on *Scope*. *Successes* succeeded, *Failures* failed, *No-Ops* no-ops. |

## COMPONENT_FIRMWARE_UPGRADE_STARTED

| Severity | Description |
|---|---|
| informational | Starting *Upgrade type* upgrade of *Firmware type* firmware, version *Label*, on *Scope*. |

## COMPONENT_FIRMWARE_CANNOT_PHASEOUT_COMPONENT

| Severity | Description |
|---|---|
| minor | Cannot phase out *Component ID*: *Error*. Firmware upgrade result was: *Upgrade result*. |

## COMPONENT_FIRMWARE_CANNOT_FAIL_COMPONENT

| Severity | Description |
|---|---|
| minor | Cannot fail *Component ID*: *Error*. Firmware upgrade result was: *Upgrade result*. |

## MIRRORING_CONNECTIVITY_TO_NON_XIV_TARGET

| Severity | Description | Troubleshooting |
|---|---|---|
| warning | Gateway Node #*Node ID*: connection to *target name*:*target's connection index* mirroring connection was established, but being ignored because the remote end is not an XIV target or is not properly configured | Please make sure the target's designation is correct, that the connection's parameters identify the intended system and that the intended system has a target_port defined for this system. |

## DM_CONNECTIVITY_TO_XIV_TARGET

| Severity | Description | Troubleshooting |
|---|---|---|
| warning | Gateway Node #*Node ID*: connection to *target name*:*target's connection index* DM connection was established, but being ignored because the remote end is an XIV target configured for mirroring, rather than a host | Please make sure the target's designation is correct, that the connection's parameters identify the intended system and that the intended system has a host defined for this system (and not a target_port). |

## EMERGENCY_ROOT_ACCESS

| Severity | Description |
|---|---|
| warning | Emergency login to 'root' account on module '*Component ID*' from '*IP Address*' using key number '*Authorized Key Number*'. |

## EMERGENCY_CONSOLE_ACCESS

| Severity | Description |
|---|---|
| warning | Emergency login to '*Unix Account Name*' account on module '*Component ID*' from tty '*TTY Device*'. |

## CR_BYPASS_ACCESS

| Severity | Description |
|---|---|
| warning | *Command that bypasses CR mechanism* access to '*Unix Account Name*' account on module '*Component ID*' from '*IP Address*'. |

## CR_KEY_SETUP_OK

| Severity | Description |
|---|---|
| informational | Challenge-response key was successfully set on all modules in the system. |

## CR_KEY_UPGRADE_NOT_DONE

| Severity | Description |
|---|---|
| warning | Challenge-response key was not upgraded on the system since a valid key has been previously set. |

## CR_KEY_SETUP_FAILED

| Severity | Description |
|---|---|
| major | Failed to set challenge-response key on module '*Component ID*'. |

## SSH_REVOKE_KEY_OK

| Severity | Description |
|---|---|
| informational | Authorized SSH key ending with '*Tail of Authorized SSH key*' was successfully revoked for user '*Unix Account Name*' on all modules in the system. |

## SSH_REVOKE_KEY_FAILED

| Severity | Description |
|---|---|
| major | Failed to revoke authorized SSH key ending with '*Tail of Authorized SSH key*' for user '*Unix Account Name*' on module '*Component ID*'. |

## IB_SWITCH_PHASEOUT_STARTED

| Severity | Description |
|---|---|
| informational | System started phasing out *Component ID*. |

## IB_SWITCH_PHASEIN_STARTED

| Severity | Description |
|---|---|
| informational | System started phasing in *Component ID*. |

## IB_SWITCH_PHASEIN_FAILED

| Severity | Description |
|---|---|
| warning | *Component ID* has failed to phase-in. |

## IB_SWITCH_CONFIG_FAILED

| Severity | Description | Troubleshooting |
|---|---|---|
| warning | *Component ID* could not be configured | Contact IBM Support |

## IB_SWITCH_FIRMWARE_INCOMPATIBLE

| Severity | Description | Troubleshooting |
|---|---|---|
| warning | The firmware version of *Component ID* is '*New Version*'. It should be '*Old Version*'. | None |

## IB_SWITCH_CPLD_INCOMPATIBLE

| Severity | Description | Troubleshooting |
|---|---|---|
| warning | The CPLD version of *Component ID* is invalid. | None |

## IB_SWITCH_FIRMWARE_UPDATE_IN_PROGRESS

| Severity | Description | Troubleshooting |
|---|---|---|
| informational | Firmware version of *Component ID* is '*Old Version*'. It should be '*New Version*'. Firmware will be updated. It may take a while. | Wait for IB switch to complete initialization. |

## IB_SWITCH_FIRMWARE_UPDATED

| Severity | Description | Troubleshooting |
|---|---|---|
| informational | The firmware version of *Component ID* was updated to '*New Version*'. | None. |

## IB_SWITCH_LOG_COLLECT_OK

| Severity | Description |
|---|---|
| informational | Log collection for IB switch '*switch_id*' completed successfuly. Log can be found in module '*log_module*' in the following directory: '*log_location*'. |

## IB_SWITCH_LOG_COLLECT_FAILED

| Severity | Description |
|---|---|
| warning | Log collection for IB switch '*switch_id*' failed. Failure reason: '*failure_reason*'. |

## IB_SWITCH_MGMT_LINK_AVAILABLE

| Severity | Description |
|---|---|
| informational | Management link *Type* of *Component ID* is available. |

## IB_SWITCH_MGMT_LINK_UNAVAIL

| Severity | Description |
|---|---|
| warning | Management link *Type* of *Component ID* is unavailable. |

## IB_SWITCH_MGMT_LINK_MISWIRED

| Severity | Description | Troubleshooting |
|---|---|---|
| warning | Management link *Type* of *Component ID* is connected to *Router* instead of *Expected Router*. | Check wiring |

## IB_SWITCH_MGMT_AVAILABLE

| Severity | Description |
|---|---|
| informational | Management of *Component ID* is available. |

## IB_SWITCH_MGMT_UNAVAIL

| Severity | Description |
|---|---|
| major | Management of *Component ID* is unavailable. |

## IB_SWITCH_PSU_OK

| Severity | Description |
|---|---|
| informational | *Component ID* has returned to normal state. |

## IB_SWITCH_PSU_MONITOR_FAILED

| Severity | Description |
|---|---|
| minor | *Component ID* sensor cannot be read. |

## IB_SWITCH_PSU_IS_MISSING

| Severity | Description |
|---|---|
| minor | *Component ID* is not present. |

## IB_SWITCH_PSU_FAIL

| Severity | Description |
|---|---|
| major | *Component ID* failed. |

## IB_SWITCH_BBU_OK

| Severity | Description |
|---|---|
| informational | *Component ID* has returned to normal state. |

## IB_SWITCH_BBU_MONITOR_FAILED

| Severity | Description |
|---|---|
| minor | *Component ID* sensor cannot be read. |

## IB_SWITCH_BBU_IS_MISSING

| Severity | Description |
|---|---|
| minor | *Component ID* is not present. |

## IB_SWITCH_BBU_FAIL

| Severity | Description |
|---|---|
| major | *Component ID* failed. |

## IB_SWITCH_BBU_END_OF_LIFE

| Severity | Description |
|---|---|
| major | *Component ID* has reached end of life. |

## IB_SWITCH_FAN_OK

| Severity | Description |
|---|---|
| informational | *Component ID* has returned to normal state. |

## IB_SWITCH_FAN_MONITOR_FAILED

| Severity | Description |
|---|---|
| minor | *Component ID* sensor cannot be read. |

## IB_SWITCH_FAN_IS_MISSING

| Severity | Description |
|---|---|
| minor | *Component ID* is not present. |

## IB_SWITCH_FAN_FAIL

| Severity | Description |
|---|---|
| major | *Component ID* failed. |

## IB_SWITCH_PSU_FAN_OK

| Severity | Description |
|---|---|
| informational | *Component ID* has returned to normal state. |

## IB_SWITCH_PSU_FAN_MONITOR_FAILED

| Severity | Description |
|---|---|
| minor | *Component ID* sensor cannot be read. |

## IB_SWITCH_PSU_FAN_IS_MISSING

| Severity | Description |
|---|---|
| minor | *Component ID* is not present. |

## IB_SWITCH_PSU_FAN_FAIL

| Severity | Description |
|---|---|
| major | *Component ID* failed. |

## IB_SWITCH_VOLTAGE_MONITOR_FAILED

| Severity | Description |
|---|---|
| minor | *Component ID* voltage sensor *Sensor Type* cannot be read. |

## IB_SWITCH_VOLTAGE_CHANGE

| Severity | Description |
|---|---|
| informational | *Component ID* voltage sensor *Sensor Type* changed from *Old Status* to *Status*. |

## IB_SWITCH_PSU_VOLTAGE_MONITOR_FAILED

| Severity | Description |
|---|---|
| minor | *Component ID* voltage sensor cannot be read. |

## IB_SWITCH_PSU_VOLTAGE_CHANGE

| Severity | Description |
|---|---|
| informational | *Component ID* voltage sensor changed from *Old Status* to *Status*. |

## IB_SWITCH_BBU_VOLTAGE_MONITOR_FAILED

| Severity | Description |
|---|---|
| minor | *Component ID* voltage sensor cannot be read. |

## IB_SWITCH_BBU_VOLTAGE_CHANGE

| Severity | Description |
|---|---|
| informational | *Component ID* voltage sensor changed from *Old Status* to *Status*. |

## IB_SWITCH_TEMPERATURE_MONITOR_FAILED

| Severity | Description |
|---|---|
| minor | *Component ID* sensor *Sensor Type* temperature cannot be read. |

## IB_SWITCH_TEMPERATURE_OK

| Severity | Description |
|---|---|
| informational | *Component ID* sensor *Sensor Type* temperature has returned to normal state. |

## IB_SWITCH_TEMPERATURE_HIGH

| Severity | Description |
|---|---|
| minor | *Component ID* sensor *Sensor Type* temperature is high. |

## IB_SWITCH_TEMPERATURE_CRITICAL

| Severity | Description |
|---|---|
| major | *Component ID* sensor *Sensor Type* temperature is critical. |

## IB_SWITCH_PSU_TEMPERATURE_MONITOR_FAILED

| Severity | Description |
|---|---|
| minor | *Component ID* temperature cannot be read. |

## IB_SWITCH_PSU_TEMPERATURE_OK

| Severity | Description |
|---|---|
| informational | *Component ID* temperature has returned to normal state. |

## IB_SWITCH_PSU_TEMPERATURE_HIGH

| Severity | Description |
|---|---|
| minor | *Component ID* temperature is high. |

## IB_SWITCH_PSU_TEMPERATURE_CRITICAL

| Severity | Description |
|---|---|
| major | *Component ID* temperature is critical. |

## IB_SWITCH_BBU_TEMPERATURE_MONITOR_FAILED

| Severity | Description |
|---|---|
| minor | *Component ID* temperature cannot be read. |

## IB_SWITCH_BBU_TEMPERATURE_OK

| Severity | Description |
|---|---|
| informational | *Component ID* temperature has returned to normal state. |

## IB_SWITCH_BBU_TEMPERATURE_HIGH

| Severity | Description |
|---|---|
| minor | *Component ID* temperature is high. |

## IB_SWITCH_BBU_TEMPERATURE_CRITICAL

| Severity | Description |
|---|---|
| major | *Component ID* temperature is critical. |

## IB_PORT_MOVED

| Severity | Description |
|---|---|
| informational | Infinibind module port '*module_port*' moved from '*from_port*' to '*to_port*'. |

## SYSTEM_TEMPERATURE_IS_ABOVE_NORMAL

| Severity | Description | Troubleshooting |
|---|---|---|
| warning | System temperature is *System Temperature*C, which is above the normal temperature. | Cool the system down. |

## SYSTEM_TEMPERATURE_IS_HIGH

| Severity | Description | Troubleshooting |
|---|---|---|
| minor | System temperature is *System Temperature*C, which is high. | Cool the system down. |

## SYSTEM_TEMPERATURE_IS_CRITICALLY_HIGH

| Severity | Description | Troubleshooting |
|---|---|---|
| critical | System temperature is *System Temperature*C, which exceeds operational level. Please initiate shutdown sequence. Without further action, the system will automatically shut itself down if it reaches *Shutdown Threshold*C. | Cool the system down immediately or shut down the system using 'shutdown -y' and contact IBM support. |

## SYSTEM_TEMPERATURE_IS_CRITICALLY_HIGH_SHUTTING_DOWN

| Severity | Description | Troubleshooting |
|---|---|---|
| critical | System temperature is *System Temperature*C, which is critically high. Shutting down the system. | Shut down the system using 'shutdown -y' and contact IBM support. |

## SYSTEM_TEMPERATURE_IS_TOO_HIGH

| Severity | Description | Troubleshooting |
|---|---|---|
| major | System temperature is *System Temperature*C. It approaches the maximal allowable value. | Cool the system down and contact IBM support. |

## SYSTEM_TEMPERATURE_IS_TOO_LOW

| Severity | Description | Troubleshooting |
|---|---|---|
| major | System temperature is *System Temperature*C, which is lower than the minimal allowable value. | Contact IBM Support |

## SYSTEM_TEMPERATURE_IS_OK_NOW

| Severity | Description |
|---|---|
| informational | System temperature is *System Temperature*C, which is within allowed limits. |

## SYSTEM_AVERAGE_POWER_PREPARATION_STARTED

| Severity | Description |
|---|---|
| informational | System average power consumption preparation has started, when it's over the up-to-date value will be available. |

## SYSTEM_AVERAGE_POWER_PREPARATION_OVER

| Severity | Description |
|---|---|
| informational | System average power consumption preparation is over, you can now read the up-to-date value. |

## SYSTEM_AVERAGE_POWER_PREPARATION_FAILED

| Severity | Description |
|---|---|
| warning | System average power consumption preparation has failed, try again later. |

## ENCRYPT_ENABLE_DRIVE_FAILED

| Severity | Description | Troubleshooting |
|---|---|---|
| major | Failed to enable encryption for *Component ID*. Error code: *Failure Reason*. | Contact IBM Support |

## ENCRYPT_ENABLE_VAULT_DEVICE_FAILED

| Severity | Description | Troubleshooting |
|---|---|---|
| major | Failed to enable encryption for *Component ID*. Error code: *Failure Reason*. | Contact IBM Support |

## VAULT_DEVICE_ENCRYPTING_ENABLE_FAILED

| Severity | Description | Troubleshooting |
|---|---|---|
| major | Failed to enable encryption for *Component ID*. Error code: *Failure Reason*. | Contact IBM Support |

## VAULT_DEVICE_SECURE_ERASE_FAILED

| Severity | Description | Troubleshooting |
|---|---|---|
| major | Failed to secure erase *Component ID*. Error code: *Failure Reason*. | Contact IBM Support |

## VAULT_DEVICE_SECURE_ERASE_SUCCESSFUL

| Severity | Description | Troubleshooting |
|---|---|---|
| informational | Secure erase was successful for *Component ID*. | Contact IBM Support |

## VAULT_DEVICE_SECURE_ERASE_NOT_DONE

| Severity | Description | Troubleshooting |
|---|---|---|
| informational | Secure erase was not done for *Component ID* as it is unsecured. | Contact IBM Support |

## MODULE_BBU_OVERHEATING

| Severity | Description | Troubleshooting |
|---|---|---|
| major | *bbu* reached a temperature of *temp*C, above *critical_set*C. BBU will be disconnected until it cools down below *critical_release*C. | Contact IBM Support |

## MODULE_BBU_IS_DISCHARGING

| Severity | Description |
|---|---|
| informational | *BBU id* changed state from '*old_state*' to '*new state*'. |

## MODULE_BBU_STOPPED_DISCHARGING

| Severity | Description |
|---|---|
| informational | *BBU id* changed state from '*old_state*' to '*new state*'. |

## MODULE_BBU_CHARGING_WAS_EXPLICITLY_ENABLED

| Severity | Description |
|---|---|
| minor | *BBU id* was not charging, it had to be reset explicitly. |

## MODULE_BBU_NOT_CHARGING_AFTER_RESET

| Severity | Description |
|---|---|
| major | *BBU id* is still not charging after *Reset Attempts* reset attempts. |

## MODULE_BBU_NOT_CHARGED_AFTER_RESET

| Severity | Description |
|---|---|
| major | *BBU id* was reset and it is charging, but after *Minutes* minutes it is still only *Percent Charged*% charged, which is not enough. |

## MODULE_BBU_STILL_NOT_SUFFICIENTLY_CHARGED

| Severity | Description |
|---|---|
| major | *BBU id* is charging, but after *Minutes* minutes it is still only *Percent Charged*% charged, which is not enough. |

## MODULE_BBU_DISCHARGING_WAS_EXPLICITLY_ENABLED

| Severity | Description |
|---|---|
| minor | *BBU id* was in discharge disabled mode, it had to be enabled explicitly. |

## MODULE_BBU_IS_CHARGING

| Severity | Description |
|---|---|
| informational | *BBU id* changed state from '*old_state*' to '*new state*'. |

## MODULE_BBU_IS_FULL

| Severity | Description |
|---|---|
| informational | *BBU id* changed state from '*old_state*' to '*new state*'. |

## MODULE_BBU_DRIVER_NOT_LOADED

| Severity | Description | Troubleshooting |
|---|---|---|
| major | BBU driver is not loaded in module *Component*. modprobe ruby | Contact IBM Support |

## MODULE_BBU_CONTROLLER_NOT_PRESENT

| Severity | Description | Troubleshooting |
|---|---|---|
| major | BBU controller board not detected in module *Component*. | Contact IBM Support |

## MODULE_IS_MISSING_EPOW_CABLE

| Severity | Description | Troubleshooting |
|---|---|---|
| major | The EPOW cable in the BBU controller board on module *Component* is not detected. | Contact IBM Support |

## MODULE_EPOW_CABLE_OK_NOW

| Severity | Description |
|---|---|
| informational | The EPOW cable in the BBU controller board on module *Component* is now OK. |

## MODULE_IS_MISSING_POWER_SENSE_CABLE

| Severity | Description | Troubleshooting |
|---|---|---|
| major | The power sense cable in the BBU controller board on module *Component* is not detected. | Contact IBM Support |

## MODULE_POWER_SENSE_CABLE_OK_NOW

| Severity | Description |
|---|---|
| informational | The power sense cable in the BBU controller board on module *Component* is now OK. |

## MODULE_BBU_CALIBRATION_STARTED

| Severity | Description |
|---|---|
| informational | *BBU id* started calibration. |

## MODULE_BBU_CALIBRATION_ENDED

| Severity | Description |
|---|---|
| informational | *BBU id* ended calibration with status '*result*'. |

## MODULE_BBU_TEMPERATURE_TOO_HIGH_FOR_CALIBRATION

| Severity | Description | Troubleshooting |
|---|---|---|
| major | *BBU id* temperature is *Temperature.Temperature Tenths*C which is too high. Calibration is stopped. | Contact IBM Support |

## MODULE_BBU_TEST_IN_CHARGING_MODE

| Severity | Description |
|---|---|
| informational | Test of *Component ID* is pending and will resume once it gets to a capacity of *Target Capacity%*. The current capacity is *Current Capacity%*. |

## SDR_PSU_STATUS_OK

| Severity | Description |
|---|---|
| informational | *Psu* (location *Location*) is now OK. Changed from '*previous_sdr_status*' to '*sdr_status*'. |

## SDR_PSU_STATUS_BAD

| Severity | Description |
|---|---|
| warning | *PSU* (location *Location*) is failed or off. Changed from '*previous_sdr_status*' to '*current_sdr_status*'. |

## INVALID_PSU_PART_NUMBER

| Severity | Description | Troubleshooting |
|---|---|---|
| major | *PSU* has an invalid part number '*PN*'. | Please contact IBM support and have the PSU replaced. |

## CMOS_BATTERY_TOO_WEAK

| Severity | Description | Troubleshooting |
|---|---|---|
| major | The CMOS battery on *Module* is too weak. | Please contact IBM support and have the battery replaced. |

## CMOS_BATTERY_IS_OK

| Severity | Description |
|---|---|
| informational | The CMOS battery on *Module* is now OK. |

## FC_LINK_IS_NOW_DOWN

| Severity | Description | Troubleshooting |
|---|---|---|
| major | FC port *Component* Active Firmware *Firmware version* - link disconnected. | Contact IBM Support |

## FC_LINK_IS_NOW_UP

| Severity | Description | Troubleshooting |
|---|---|---|
| informational | FC port *Component* - link regained. | Contact IBM Support |

## FC_LINK_SYNC_ERROR

| Severity | Description | Troubleshooting |
|---|---|---|
| major | FC port *Component* - errors on the physical layer: *Reason*. | Please contact support. |

## FC_PORT_TEST_STARTED

| Severity | Description |
|---|---|
| informational | FC port *Component* - test started |

## FC_PORT_TEST_NOT_STARTED

| Severity | Description | Troubleshooting |
|---|---|---|
| informational | FC port *Component* - test not started. | Check port state |

## FC_PORT_TEST_FAILED

| Severity | Description | Troubleshooting |
|---|---|---|
| major | FC port *Component* - test failed. | Contact IBM Support |

## FC_PORT_TEST_SUCCESS

| Severity | Description |
|---|---|
| informational | FC port *Component* - test success. |

## FC_PORT_TEST_ABORTED

| Severity | Description |
|---|---|
| informational | FC port *Component* - test aborted. |

## COMPONENT_NETWORK_LINK_IS_DOWN

| Severity | Description | Troubleshooting |
|---|---|---|
| major | Network interface to *Connected Component* on *Component ID* - link disconnected. | Contact IBM Support |

## COMPONENT_NETWORK_LINK_IS_UP

| Severity | Description | Troubleshooting |
|---|---|---|
| informational | Network interface to component *Connected Component* on *Component ID* - link regained. | Contact IBM Support |

## MODULE_IS_MISSING_REQUIRED_MEMORY

| Severity | Description | Troubleshooting |
|---|---|---|
| major | *Component ID* has less memory (*actual_mem* GB) than is defined for use (*req_mem* GB). | Please contact your Administrator. |

## POD_IB_MISWIRE

| Severity | Description | Troubleshooting |
|---|---|---|
| warning | POD module miswired: *Module ID*. | Contact IBM Support |

## POD_IB_MISWIRE_CORRECTED

| Severity | Description |
|---|---|
| informational | POD module miswire corrected: *Module ID*. |

## MODULE_SET_LED_LOCATOR_FAILED

| Severity | Description |
|---|---|
| warning | Failed to set LED locator on *module*. |

## MODULE_SET_LED_LOCATOR_COMPLETED

| Severity | Description |
|---|---|
| informational | LED locator set successfully on *module*. |

## PERF_CLASS_RESOURCE_EXHAUSTION

| Severity | Description |
|---|---|
| warning | Exhausted all allowed resources for performance classes on *Module Id*, BUSY until resources available. |

## CONNECTED_HOSTS_LIMIT_REACHED

| Severity | Description |
|---|---|
| informational | Number of connected Hosts was reached for port '*port_id*' in Module *Module Id*. |

## QoS_HAS_BEEN_TRIGGERED

| Severity | Description |
|---|---|
| informational | Queues on port '*port_id*' in Module *Module Id* caused QoS to be activated. |

## PERF_CLASS_RATE_AT_LIMIT

| Severity | Description |
|---|---|
| informational | Performance class '*perf_class*' on *Module Id* reached its limit of *Limit Limit Name*, IOs being throttled. |

## INDEPENDENT_PERF_CLASS_RATE_AT_LIMIT

| Severity | Description |
|---|---|
| informational | Performance class '*perf_class*' object *type:name* on *Module Id* reached its limit of *Limit Limit Name*, IOs being throttled. |

## PORT_PREP_FOR_UPGRADE_TIMED_OUT

| Severity | Description |
|---|---|
| warning | Preparation of *port_type* port '*local_port_name*' for hot-upgrade timed out due to host '*host_name*' port '*host_port_name*'*host_port_addr* |

## INTERFACE_DISCONNECTED_FROM_TARGET

| Severity | Description |
|---|---|
| major | Interface node on module *module* cannot access target '*target*' through any gateway module. |

## INTERFACE_RECONNECTED_TO_TARGET

| Severity | Description |
|---|---|
| major | Interface node on module *module* can access target '*target*'. |

## METADATA_SERVICE_DB_CREATE

| Severity | Description |
|---|---|
| informational | Database *DB* was created |

## METADATA_SERVICE_DB_DELETE

| Severity | Description |
|---|---|
| informational | Database *DB* was deleted |

## IPINTERFACE_CREATE

| Severity | Description |
|---|---|
| informational | A new iSCSI IP interface was defined with name '*interface name*' on module *module* with port '*port list*' and IP address *IP address*. |

## IPINTERFACE_DELETE

| Severity | Description |
|---|---|
| informational | iSCSI IP interface with name '*interface name*' was deleted. |

## IPINTERFACE_RENAME

| Severity | Description |
|---|---|
| informational | iSCSI IP interface with name '*old name*' and was renamed '*interface name*'. |

## IPINTERFACE_UPDATE

| Severity | Description |
|---|---|
| informational | ISCSI IP interface with name '*interface name*' was updated. Its IP address is *IP address*. |

## IPINTERFACE_UPDATE_MANAGEMENT

| Severity | Description |
|---|---|
| informational | Management IP interfaces were updated. Management IPs are *IP addresses*. |

## IPINTERFACE_UPDATE_MANAGEMENT_IPV6

| Severity | Description |
|---|---|
| informational | Management IP interfaces were updated. Management IPv6 addresses are *IPv6 addresses*. |

## IPINTERFACE_UPDATE_VPN

| Severity | Description |
|---|---|
| informational | VPN IP interfaces were updated. VPN IPs are *IP addresses*. |

## IPINTERFACE_VLAN_UPDATE

| Severity | Description |
|---|---|
| informational | System VLAN configuration was updated: PCP=*PCP* |

## IPINTERFACE_UPDATE_VPN_IPV6

| Severity | Description |
|---|---|
| informational | VPN IPv6 interfaces were updated. VPN IPv6 addresses are *IP addresses*. |

## ETHERNET_PORT_UPDATE

| Severity | Description |
|---|---|
| informational | Ethernet port with name *Ethernet port* was updated. |

## AUXILIARY_INTERNAL_PORTS_ENABLED

| Severity | Description |
|---|---|
| informational | *Port Count* auxiliary internal Ethernet ports were enabled. |

## AUXILIARY_INTERNAL_PORTS_DISABLED

| Severity | Description |
|---|---|
| informational | *Port Count* auxiliary internal Ethernet ports were disabled. |

## IPSEC_ENABLED

| Severity | Description |
|---|---|
| informational | IPSec was enabled |

## IPSEC_DISABLED

| Severity | Description |
|---|---|
| informational | IPSec was disabled |

## IPSEC_CONNECTION_ADDED

| Severity | Description |
|---|---|
| informational | A new IPSec connection named '*name*' was added |

## IPSEC_CONNECTION_UPDATED

| Severity | Description |
|---|---|
| informational | The IPSec connection named '*name*' was updated |

## IPSEC_CONNECTION_REMOVED

| Severity | Description |
|---|---|
| informational | The IPSec connection named '*name*' was removed |

## PRIVATE_KEY_ADDED

| Severity | Description |
|---|---|
| informational | A new private key named '*name*' with fingerprint '*fingerprint*' and size *key_size* bits was added. |

## CERTIFICATE_REMOVED

| Severity | Description |
|---|---|
| informational | The certificate named '*name*' was removed. |

## PKCS12_CERTIFICATE_ADDED

| Severity | Description |
|---|---|
| informational | A new PKCS#12 named '*name*' with fingerprint '*fingerprint*' was added. |

## PKI_RENAME

| Severity | Description |
|---|---|
| informational | PKI with the name '*old name*' was renamed to '*new name*' |

## PKI_UPDATED

| Severity | Description |
|---|---|
| informational | PKI with the name '*name*' and fingerprint '*fingerprint*' was updated |

## PROTOCOL_CONFIGURATION_CHANGED

| Severity | Description |
|---|---|
| informational | Protocol configuration changed for protocol '*Protocol Type*'. |

## EMAIL_HAS_FAILED

| Severity | Description | Troubleshooting |
|---|---|---|
| variable | Sending event *Event Code* (*Event Index*) to *Destination List* via *SMTP Gateway* failed. Module: *Module ID*; Error message: '*Error Message*'; timeout expired: *Timeout Expired?*. | Contact IBM Support |

## BULK_EMAIL_HAS_FAILED

| Severity | Description | Troubleshooting |
|---|---|---|
| variable | Sending bulk email with *Events Number* events to *Destination List* via *SMTP Gateway* failed. Module: *Module ID*; Error message: '*Error Message*'; timeout expired: *Timeout Expired?*. | Contact IBM Support |

## SMS_HAS_FAILED

| Severity | Description | Troubleshooting |
|---|---|---|
| variable | Sending event *Event Code* (*Event Index*) to *Destination List* via *SMS Gateway* and *SMTP Gateway* failed. Module: *Module ID*; Error message: '*Error Message*'; timeout expired: *Timeout Expired?*. | Contact IBM Support |

## HTTPS_HAS_FAILED

| Severity | Description | Troubleshooting |
|---|---|---|
| variable | Sending event *Event Code* (*Event Index*) to *Destination List* via *HTTPS address* failed. Module: *Module ID*; Error message: '*Error Message*' (*HTTP error code*); timeout expired: *Timeout Expired?*. | Contact IBM Support |

## EMAIL_NOT_SENT

| Severity | Description | Troubleshooting |
|---|---|---|
| variable | Sending event *Event Code* (*Event Index*) to *Destination List* via *SMTP Gateway* was waived because of failed SMTP gateway. It will be not be used until *Retry Time*. | Contact IBM Support |

## SMS_NOT_SENT

| Severity | Description | Troubleshooting |
|---|---|---|
| variable | Sending event *Event Code* (*Event Index*) to *Destination List* via *SMS Gateway* and *SMTP Gateway* was waived because of failed SMTP gateway. It will be not be used until *Retry Time*. | Contact IBM Support |

## HEARTBEAT_EMAIL_HAS_FAILED

| Severity | Description | Troubleshooting |
|---|---|---|
| minor | Sending heartbeat to *Destination Name* via *SMTP Gateway* failed. Module: *Module ID*; Error message: '*Error Message*'; timeout expired: *Timeout Expired?*. | Contact IBM Support |

## HEARTBEAT_SMS_HAS_FAILED

| Severity | Description | Troubleshooting |
|---|---|---|
| minor | Sending heartbeat to *Destination Name* via *SMS Gateway* and *SMTP Gateway* failed. Module: *Module ID*; Error message: '*Error Message*'; timeout expired: *Timeout Expired?*. | Contact IBM Support |

## TEST_EMAIL_HAS_FAILED

| Severity | Description | Troubleshooting |
|---|---|---|
| minor | Sending test to *Destination Name* via *SMTP Gateway* failed. Module: *Module ID*; Error message: '*Error Message*'; timeout expired: *Timeout Expired?*. | Contact IBM Support |

## TEST_SMS_HAS_FAILED

| Severity | Description | Troubleshooting |
|---|---|---|
| minor | Sending test to *Destination Name* via *SMS Gateway* and *SMTP Gateway* failed. Module: *Module ID*; Error message: '*Error Message*'; timeout expired: *Timeout Expired?*. | Contact IBM Support |

## CUSTOM_EVENT

| Severity | Description |
|---|---|
| variable | *Description* |

## UPGRADE_SOFTWARE_DOWNLOAD_FINISHED

| Severity | Description |
|---|---|
| informational | Finished downloading software needed for upgrade to version *version*. Upgrade consequence is *consequence* |

## UPGRADE_FILE_LIST_RETRIEVAL_FAILED

| Severity | Description | Troubleshooting |
|---|---|---|
| critical | Could not receive new version's file list from repository. Error code is *error*. | Contact IBM Support |

## UPGRADE_STARTS

| Severity | Description |
|---|---|
| informational | System starting an upgrade. |

## PRE_UPGRADE

| Severity | Description |
|---|---|
| informational | System preparing an upgrade procedure type *type* . |

## UPGRADE_IS_OVER

| Severity | Description |
|---|---|
| informational | System went up after an upgrade. |

## IOS_RESTORED_AFTER_HOT_UPGRADE

| Severity | Description |
|---|---|
| informational | System is able to perform I/Os after a hot upgrade. |

## UPGRADE_NO_NEW_FILES_FOR_UPGRADE

| Severity | Description | Troubleshooting |
|---|---|---|
| warning | Repository version does not contain any new files. current version *current_version* new version is *new_version* | Contact IBM Support |

## UPGRADE_DOWNLOAD_REPOSITORY_COPY

| Severity | Description | Troubleshooting |
|---|---|---|
| critical | Mirroring needed files from repository failed. Mirroring module is *mirroring_module* error is *error* | Contact IBM Support |

## UPGRADE_LOCAL_VERSION_DOWNLOAD_FAILED

| Severity | Description | Troubleshooting |
|---|---|---|
| critical | Failure to distribute new sofware internally. Error code is *error* . | Contact IBM Support |

## UPGRADE_WAS_CANCELLED

| Severity | Description | Troubleshooting |
|---|---|---|
| informational | Upgrade was cancelled with reason *reason* . | Contact IBM Support |

## HOT_UPGRADE_ABORTED

| Severity | Description | Troubleshooting |
|---|---|---|
| critical | Hot upgrade aborted with reason *reason*. | Contact IBM Support |

## HOT_UPGRADE_HAS_FAILED

| Severity | Description | Troubleshooting |
|---|---|---|
| critical | Hot upgrade failed while *errorneous_state*. | Contact IBM Support |

## PRE_UPGRADE_SCRIPT_INVOCATION_FAILED

| Severity | Description | Troubleshooting |
|---|---|---|
| critical | Invocation of pre-upgrade script failed with error *error*. | Contact IBM Support |

## POST_UPGRADE_SCRIPT_INVOCATION_FAILED

| Severity | Description | Troubleshooting |
|---|---|---|
| critical | Invocation of post-upgrade script failed with error *error*. | Contact IBM Support |

## UPGRADE_IS_NOT_ALLOWED

| Severity | Description | Troubleshooting |
|---|---|---|
| critical | One or more of the pre-upgrade validations failed. | Fix the problems pointed out it previous events and revalidate. |

## PRE_UPGRADE_VALIDATION_FAILED

| Severity | Description | Troubleshooting |
|---|---|---|
| critical | One of the pre-upgrade validations failed with status *error*. | Contact IBM Support |

## POST_UPGRADE_SCRIPT_STARTED

| Severity | Description |
|---|---|
| informational | Post-upgrade script started. |

## POST_UPGRADE_SCRIPT_FINISHED

| Severity | Description |
|---|---|
| informational | Post-upgrade script finished successfully. |

## PRE_UPGRADE_SCRIPT_DISAPPROVES

| Severity | Description | Troubleshooting |
|---|---|---|
| critical | Upgrade cannot commence because some of the validations in the pre-upgrade script failed. Explanation: *explanation*. | Correct the system state according to the explanation and try again |

## POST_UPGRADE_SCRIPT_REPORTED_FAILURE

| Severity | Description | Troubleshooting |
|---|---|---|
| critical | Post upgrade script reported failure. Script output: *explanation*. | Correct the system state according to the explanation and try again |

## POWER_PROBLEM_CAUSING_MAINTENANCE_MODE

| Severity | Description |
|---|---|
| warning | Power state causing system to enter maintenance mode. |

## SYSTEM_ENTERED_CHARGING_STATE

| Severity | Description |
|---|---|
| informational | System cannot start work until it is sufficiently charged. |

## POWER_PROBLEM_CAUSING_MODULE_PHASEOUT

| Severity | Description |
|---|---|
| critical | Module *module* was phased out due to a power problem. |

## POWER_REPORT_PROBLEM_CAUSING_MODULE_PHASEOUT

| Severity | Description |
|---|---|
| critical | Module *module* did not report power status on time and hence it was phased out. |

## POWER_PROBLEM_CAUSING_SYSTEM_SHUTDOWN

| Severity | Description |
|---|---|
| critical | Power state causing system to shutdown due to: *Power Emergency Shutdown Reason*. |

## DELAYING_BACKUP_POWER_FAILURE_HANDLING

| Severity | Description |
|---|---|
| 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 |
|---|---|---|
| major | Switch *ib_switch* is missing. | Contact IBM Support |

## IB_SWITCH_LOST

| Severity | Description | Troubleshooting |
|---|---|---|
| warning | Switch *ib_switch* that was missing is considered lost. | Contact IBM Support |

## IB_MISSING_SWITCH_FOUND

| Severity | Description |
|---|---|
| informational | A previously missing switch *ib_switch* is now found. |

## IB_CONFIGURE_COMMAND_ERROR

| Severity | Description |
|---|---|
| warning | Infiniband configuration command *command* has failed on *component*. |

## POD_IB_PORT_MISWIRE

| Severity | Description | Troubleshooting |
|---|---|---|
| warning | POD module port miswired: *module_port* shall connect to *expected_component* but connects to *component* with GUID *guid*. | Please contact support. |

## POD_IB_PORT_MISWIRE_CORRECTED

| Severity | Description |
|---|---|
| informational | POD module port *module_port* connected to *component* miswire corrected. |

## IB_SWITCH_REBOOT_DETECTED

| Severity | Description | Troubleshooting |
|---|---|---|
| warning | IB switch '*switch_id*' has rebooted. | Contact IBM Support |

## IB_CONNECTION_SERVICES_UNAVAILABLE

| Severity | Description | Troubleshooting |
|---|---|---|
| warning | Connection services unavailable on port '*port*'. | Contact IBM Support |

## IB_CONNECTION_SERVICES_AVAILABLE

| Severity | Description | Troubleshooting |
|---|---|---|
| informational | Connection services now available on port '*port*'. | Contact IBM Support |

## REDUCED_IB_PORT_SPEED

| Severity | Description | Troubleshooting |
|---|---|---|
| major | Component *Component ID* switched to reduced link speed of *Speed* Gbps. Expected link speed is *Expected Speed* Gbps. | Contact IBM Support |

## IB_PORT_SPEED_OK

| Severity | Description | Troubleshooting |
|---|---|---|
| informational | Component *Component ID* switched to expected link speed of *Speed* Gbps. | Confirm that link blinking does not take place. |

## PERF_CLASS_MAX_IO_RATE_UPDATED

| Severity | Description |
|---|---|
| informational | Performance Class *name* max IO rate was changed to *IO rate* IOPS |

## PERF_CLASS_MAX_BW_RATE_UPDATED

| Severity | Description |
|---|---|
| informational | Performance Class *name* max BW rate was changed to *BW rate* MB/sec |

## PERF_CLASS_CREATE

| Severity | Description |
|---|---|
| informational | Performance Class with name '*name*' was created |

## PERF_CLASS_DELETE

| Severity | Description |
|---|---|
| informational | Performance Class with name '*name*' was deleted |

## PERF_CLASS_ADD_HOST

| Severity | Description |
|---|---|
| informational | Host with name '*host_name*' was added to Performance Class with name '*name*' |

## PERF_CLASS_REMOVE_HOST

| Severity | Description |
|---|---|
| informational | Host with name '*host_name*' was removed from Performance Class with name '*name*' |

## PERF_CLASS_ADD_POOL

| Severity | Description |
|---|---|
| informational | Pool with name '*pool.name*' was added to Performance Class with name '*pool.perf_class*' |

## PERF_CLASS_REMOVE_POOL

| Severity | Description |
|---|---|
| informational | Pool with name '*pool.name*' was removed from Performance Class with name '*name*' |

## PERF_CLASS_ADD_VOLUME

| Severity | Description |
|---|---|
| informational | Volume with name '*volume.name*' was added to Performance Class with name '*volume.perf_class*' |

## PERF_CLASS_REMOVE_VOLUME

| Severity | Description |
|---|---|
| informational | Volume with name '*volume.name*' was removed from Performance Class with name '*name*' |

## PERF_CLASS_ADD_DOMAIN

| Severity | Description |
|---|---|
| informational | Domain *domain_name* was added to Performance Class *name* |

## PERF_CLASS_REMOVE_DOMAIN

| Severity | Description |
|---|---|
| informational | Domain *domain_name* was removed from Performance Class *name* |

## VOLUME_MODIFIED_DURING_IO_PAUSE

| Severity | Description | Troubleshooting |
|---|---|---|
| warning | Volume '*vol_name*' of CG '*cg_name*' was modified during Pause IO with token '*token*' | Retry after completing CG changes. |

## CONS_GROUP_MODIFIED_DURING_IO_PAUSE

| Severity | Description | Troubleshooting |
|---|---|---|
| warning | CG '*cg_name*' was modified during Pause IO with token '*token*' | Retry after completing CG changes. |

## IO_PAUSED_FOR_CONS_GROUP

| Severity | Description |
|---|---|
| Informational | Pause IO on CG with name '*cg_name*' was started with *timeout*ms timeout . Token is '*token*'. |

## IO_RESUMED_FOR_CONS_GROUP_EXPLICITLY

| Severity | Description |
|---|---|
| Informational | Pause IO on CG with name '*cg_name*' and token '*token*' was resumed by user request. |

## IO_RESUMED_FOR_CONS_GROUP_AUTOMATICALLY

| Severity | Description |
|---|---|
| Informational | Pause IO on CG with name '*cg_name*' and token '*token*' was resumed after snapgroup creation. |

## IO_RESUMED_FOR_CONS_GROUP_UPON_SYSTEM_ERROR

| Severity | Description |
|---|---|
| warning | Pause IO on CG with name '*cg_name*' and token '*token*' was resumed after system error. |

## IO_RESUMED_FOR_CONS_GROUP_UPON_TIMEOUT_EXPIRATION

| Severity | Description | Troubleshooting |
|---|---|---|
| warning | Pause IO on CG with name '*cg_name*' and token '*token*' was canceled after timeout. | Use longer timeout value or require less time for performing action. |

## ALU_CREATE

| Severity | Description |
|---|---|
| informational | ALU was defined with name '*ALU name*' associated with host '*ALU host name*' lun '*ALU lun*'. |

## ALU_DELETE

| Severity | Description |
|---|---|
| informational | ALU with name '*ALU name*' associated with host '*ALU host name*' lun '*ALU lun*' was deleted. |

## ALU_UNBOUND_ALL

| Severity | Description |
|---|---|
| informational | All SLUs of ALU with name '*ALU name*' were unbound. |

## TXN_NODE_FLASH_CONNECTION_LOST

| Severity | Description |
|---|---|
| warning | TXN node on '*module*' is disconnected from FlashSystem '*FlashSystem*'. |

## TXN_NODE_FLASH_CONNECTED

| Severity | Description |
|---|---|
| informational | TXN node on '*module*' is fully connected to FlashSystem '*FlashSystem*'. |

## POD_IB_LINK_DETECTION_LINK_PERSISTENTLY_DISCONNECTED

| Severity | Description |
|---|---|
| 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 28. Return Codes

The following table specifies the return codes.

| Return Code | Error Description |
|---|---|
| 0 | Success. |
| 1 | Command execution failed. |
| 2 | No connection to the system. |
| 3 | Password is required. |
| 4 | Password does not match system password. |
| 7 | Command not allowed from this client. |
| 8 | Bad XCLI option. |
| 9 | Internal XCLI error. |

© Copyright IBM Corp. 2016, 2022

# 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.

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*

© Copyright IBM Corp. 2016, 2022

*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.